amadeus

litepaper

v1.0 - January 2026

amadeus is a high-performance Layer 1 blockchain

designed to create, deploy, and deterministically execute AI agents without requiring users to write code.

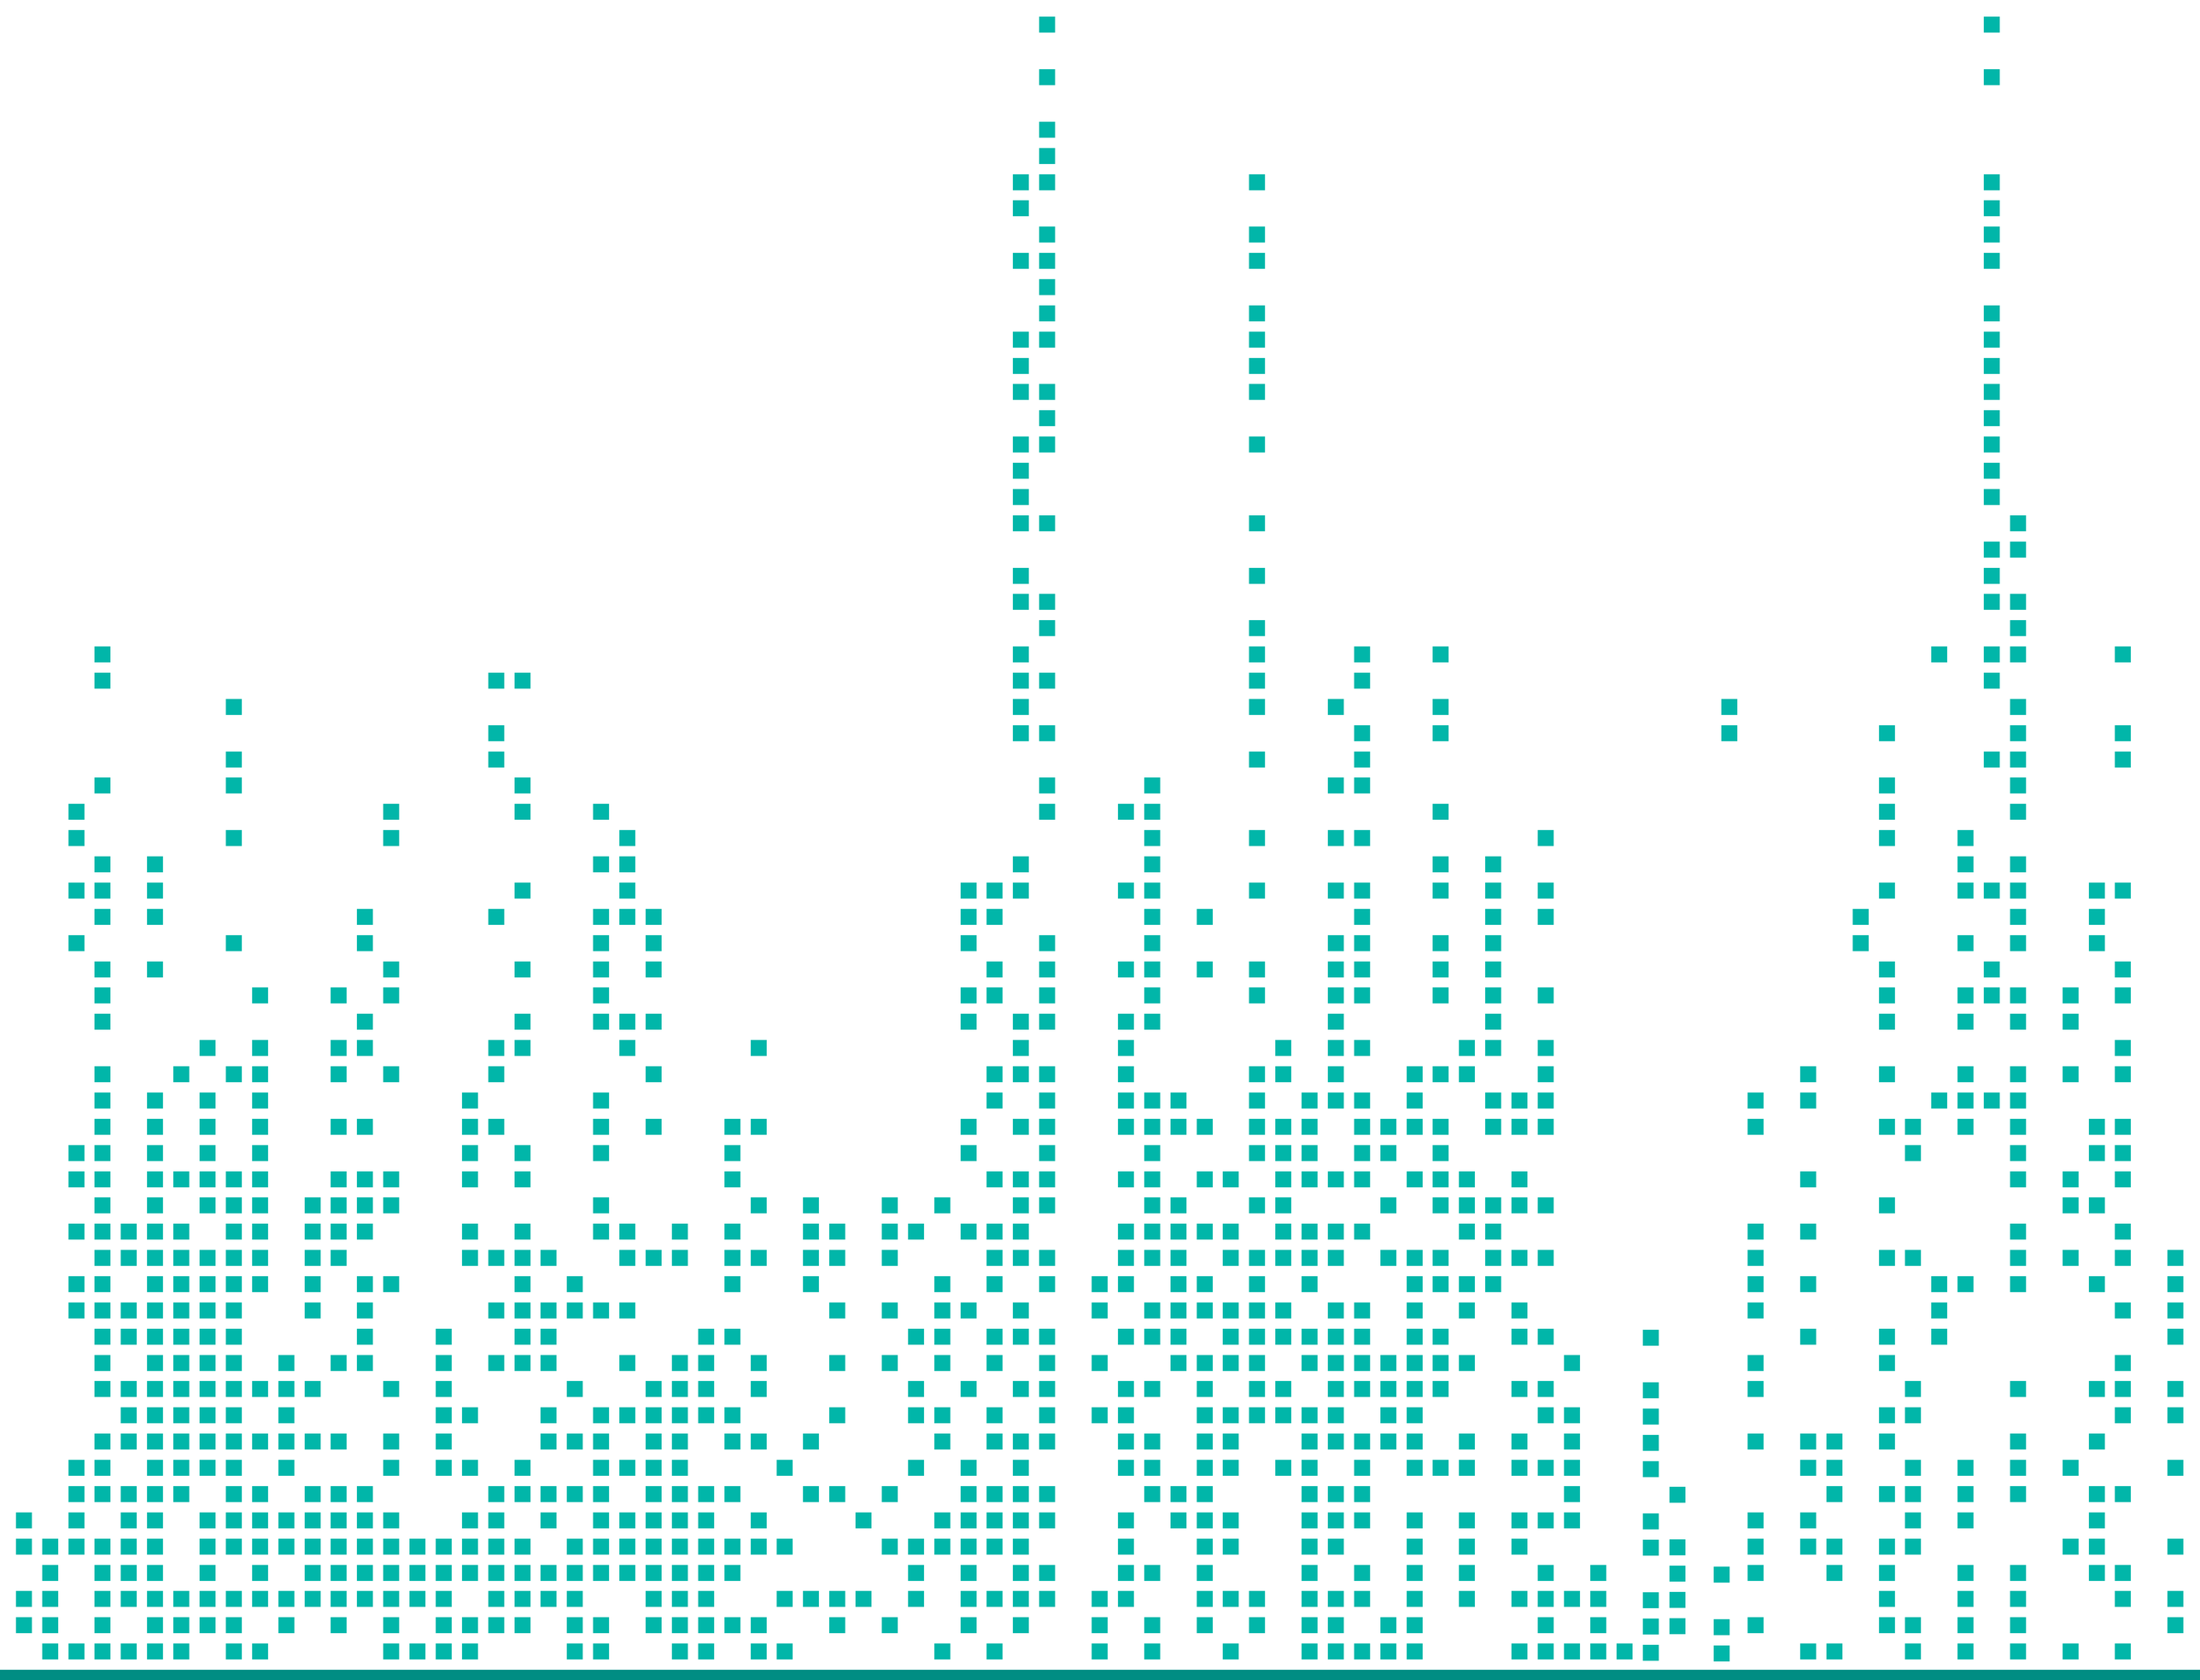

Table of Contents

Prologue

Over the next decade, most economic and informational activity will be mediated by AI agents. They will negotiate prices, route payments, rebalance portfolios, manage infrastructure, and make decisions on behalf of people, companies, and states. These agents will not be static programs: they will learn and improve over time, adapting their behavior as they observe markets, users, and other agents.

Today, that evolution happens inside opaque infrastructure. Agent logic and training data live in proprietary stacks; it is difficult to prove how an agent reached a decision, how its behavior changed over time, or whether it is still following its original constraints. At the same time, the most valuable workloads in finance and health require privacy. Existing blockchains were not designed to combine continuous agent evolution, strong privacy, and hard guarantees that an agent behaves deterministically given its inputs.

Amadeus is built to close this gap. It is a Layer 1 blockchain for private, deterministic, self-improving AI agents. Agents on Amadeus run in a deterministic execution environment: for a given version of an agent, the same inputs always produce the same outputs, and every change to that agent's logic is recorded as part of an immutable learning history. Trusted Execution Environments (TEE) provide hardware-backed privacy for sensitive workloads, so agents can evolve on private data while still being verifiable at the protocol level. The rest of this litepaper explains how Amadeus combines Useful Proof of Work, Nova AI, and TEE-backed privacy rails to create a blockchain that runs agents that think with them, predictably and provably.

Executive Summary

Amadeus is a high-performance Layer 1 blockchain built for instant asset transfers and deterministic execution of AI agents without requiring users to write code. The platform targets 0.5-second block finality, powered by a novel Useful Proof of Work (UPoW) consensus mechanism that transforms traditional mining into meaningful AI computation. At its core, Amadeus features Nova AI, a ML compiler that builds and executes agents through verifiable LLM pipelines, enabling developers to describe desired outcomes while Nova orchestrates planning, testing, and deployment. This dual-purpose architecture delivers both a high-speed general-purpose blockchain for asset transfers and a deterministic execution layer for intelligent applications where identical inputs produce identical, cryptographically verifiable outputs, creating a foundation for autonomous systems that are both provable and reproducible.

Amadeus is designed for the next generation of private, deterministic, self-improving agents:

- **Deterministic agents**: Each agent on Amadeus is compiled into a deterministic program. For a given version, its behavior is reproducible: same inputs, same outputs, anywhere on the network.

- **Self-improvement with an audit trail**: Nova AI manages the agent's lifecycle—generation, testing, deployment, and upgrades—recording each version on-chain as part of an immutable learning history. Agents can improve, but they do so in discrete, verifiable steps.

- **Privacy-preserving execution**: Upcoming TEE-backed privacy pools allow agents to run sensitive workloads—such as proprietary trading logic, private user data, factchecking, or regulated workflows—inside hardware enclaves. Their internal state remains confidential, while their deterministic behavior and evolution remain verifiable at the protocol level.

The combination of a deterministic execution, verifiable evolution and TEE based privacy, creates a neutral environment where autonomous systems can learn, transact and coordinate without exposing their internals. These capabilities are showcased in Stellaris, Amadeus' flagship vertical for private agentic trading, where deterministic, TEE-backed agents operate across DeFi, prediction markets, and agent-commerce rails. We believe it's there where deterministic, (privacy-preserving) agents can close information gaps and manage risk on behalf of users.

As of December 2025, Amadeus mainnet has processed over 205 million transactions across 44.6 million+ blocks, with 16k+ agents, 12k+ connected accounts, and 11 validators securing the network, demonstrating that the protocol is already live and used by real builders.

In December 2025 Amadeus also acquired Bitte.ai, an AI infrastructure and agent orchestration product whose technology is being integrated into Nova and Agent Studio. Bitte has already served 24,164 unique users, processed 2,849,652 messages across 344,542 chats, and registered 16,703 agents, of which 11,286 are active (≈67.6% utilization) across EVM, NEAR, and Sui, with ~93.8% of users on EVM.

The Infrastructure Gap

The AI agent market currently runs into a structural problem:

- Off-chain stacks offer powerful compute and privacy, but little verifiability: it is hard to prove what an agent actually did or how its logic has changed.

- Existing L1s and L2s provide verifiability, but were not designed for compute-heavy, continuously evolving AI workloads. They typically do not guarantee deterministic agent behavior in practice once you

Executive Summary cont.

introduce complex model calls, and they have no native mechanism to track an agent's learning history or to enforce privacy at the execution layer.

There is effectively no neutral infrastructure today where agents can:

1. Learn and improve over time,
2. Execute deterministically in real-time, and
3. Operate on private data with hardware-backed guarantees.

Amadeus addresses this by unifying a sub-second, general-purpose L1 for asset transfers with a deterministic, privacy-aware runtime for agents in a single protocol. Nova AI sits on top as the orchestration layer that turns prompts into production-ready, deterministic agents that can evolve safely.

Token and Economic Design

The AMA token serves as the network's utility instrument, designed exclusively for ecosystem operations including asset transfers, AI agent execution, transaction fee payment, miner incentivization, creator revenue sharing, and governance participation.
Amadeus launched with a fair, miner-only distribution: no VC allocation, no premine, no team lockups. 100% of the supply is emitted through UPoW. The Shenron Upgrade refined this design through a three-phase emission schedule:

- **Phase 1** - Linear Bootstrapping: Early epochs use a linear emission curve to rapidly distribute AMA and secure the network.

- **Phase 2** - Shenron Transition: At epoch 750, the network transitions from linear emissions to a smooth logarithmic decay. Total issuance per epoch drops by roughly 65%, while a simultaneous reduction in mining fees increases net miner rewards by ~43%, better aligning incentives with long-term network health.

- **Phase 3** - Logarithmic Decay & Privacy Funding: After epoch 810, emissions continue along a logarithmic decay curve toward a hard cap of 1B AMA, while a portion of rewards is reserved to fund TEE miners and privacy-preserving compute infrastructure.

This design maintains strong early incentives for miners, provides a predictable long-term supply curve, and ensures ongoing funding for the deterministic, privacy-preserving agent infrastructure that differentiates Amadeus.

With mainnet live, Amadeus already provides: sub-second finality for asset transfers, a deterministic execution environment for agents, and a production agent runtime integrated with Nova AI. The roadmap extends this foundation with: a no-code Agent Studio and marketplace, expanded Nova AI capabilities for test-driven, verifiable agent evolution, TEE-backed privacy pools and a "Privacy Shield" layer for compliant, confidential workloads, and a decentralized GPU marketplace to scale UPoW and privacy-preserving compute.

Amadeus is designed to be the settlement and coordination layer for the emerging agent-to-agent economy, where autonomous agents negotiate, trade, and collaborate on behalf of users and organizations. By combining high-performance consensus with a deterministic agent runtime, Amadeus provides the neutral, verifiable rails that allow these agents to transact and evolve safely at scale.

Overview

Amadeus is a high performance AI-native Layer 1 blockchain that runs private, deterministic, self-improving agents and gives them native payments, compute, and a marketplace without requiring users to write code.

It combines:

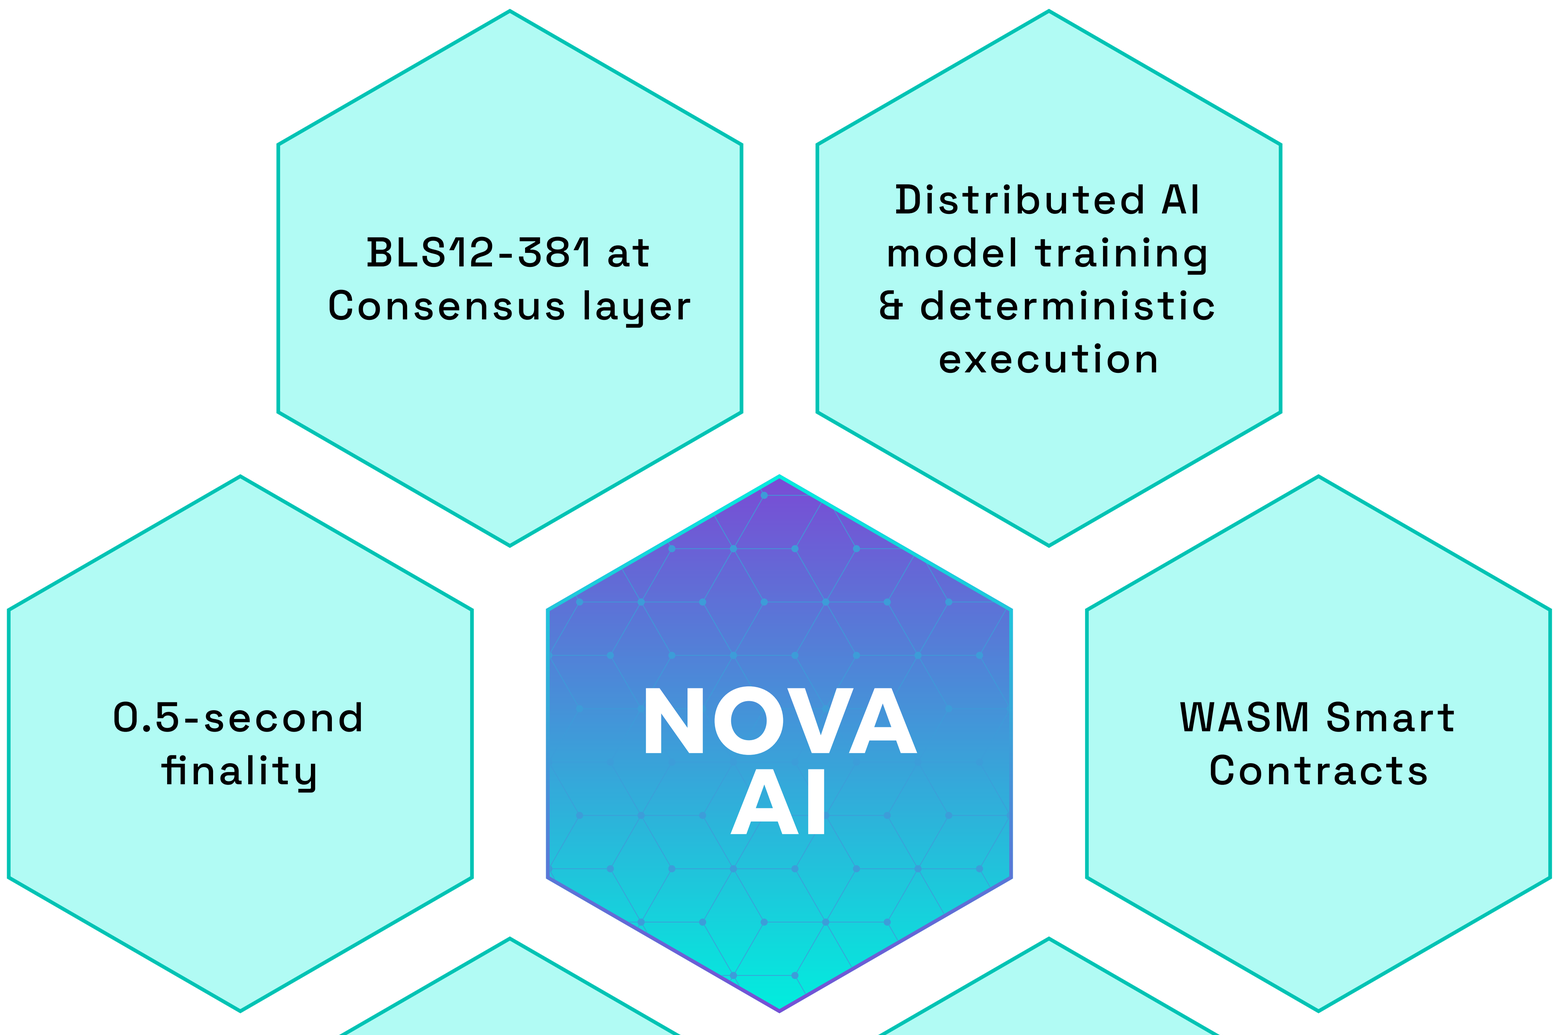

At the core of Amadeus is Nova AI, a ML orchestrator that builds and executes agents one verifiable unit at a time. Nova iterates agent logic through LLM pipelines, caching every compile, test, and design decision, enabling agents to evolve safely and coherently. Execution is deterministic: same inputs, same outputs, provable and reproducible.

From Prompt to Production Agent

Amadeus exposes its capabilities through three product pillars that turn the underlying protocol and consensus into a usable agent platform.

Agent Studio (No-Code Platform)

A visual environment for composing agents from reusable blocks. Builders drag-and-drop workflows (ingest → reason → act → verify), connect data sources, and deploy agents to Amadeus without writing low-level protocol code.

Nova AI Compiler

Developers describe goals, constraints, and policies; Nova turns these into verifiable agent logic through deterministic, test-driven LLM pipelines. Every version of an agent is compiled, tested, and recorded on-chain, so behavior is reproducible and auditable over time.

Agent Commerce & Privacy Rails

Designed to support emerging standards like x402-style HTTP payments and TEE-backed, privacy-preserving agent execution. Agents can pay for APIs, data, and compute, while selective disclosure and TEE pools allow sensitive workloads such as financial strategies, proprietary data, or regulated information to remain confidential even as they run on a public chain.

Amadeus acquires Bitte

In December 2025 Amadeus acquired Bitte.ai, an AI infrastructure and agent orchestration product. Bitte's workflow engine, integrations, and UX components are being integrated into Nova AI and Stellaris, accelerating delivery of a production-grade environment for composing, deploying, and monitoring agents on Amadeus.

At the time of acquisition, Bitte had 24,164 unique users who had exchanged 2,849,652 messages across 344,542 chats, averaging ~118 messages and ~14 chats per user. The network registered 16,703 agents, with 11,286 of them handling messages at least once (≈67.6% utilization) and 5,417 unused agents, illustrating both strong demand and a long tail of under-utilized agents. Bitte users are overwhelmingly EVM-native, with 22,669 users on EVM (93.8%), 1,432 on NEAR (5.9%), 68 on Sui (0.3%), and none on Solana yet, giving Amadeus a clear signal on where early agent demand is concentrated.

Part 1 : the problem

Wasted Computation

Traditional Proof of Work burns energy on useless hashes. Billions in hardware spend secure networks but produce no reusable compute, while AI workloads compete for the same GPUs off-chain at a premium.

Fragmented Agent Stack

Building advanced agents today requires stitching together models, APIs, schedulers, storage, and monitoring across multiple off-chain services. Existing blockchains were not designed for compute-intensive AI, forcing teams to balance scalability against verifiability. There are no truly native, no-code options for deploying and evolving agents on-chain.

No Verifiable Improvement

Agent behavior changes over time, but there is rarely a transparent, immutable record of how and why an agent evolved. Enterprises and regulators cannot trust black-box systems whose logic and training history live in proprietary backends.

Current Landscape: AI Agents Are Hard to Build

No Private, Auditable Evolution

Today's agents either run in closed infrastructure where their learning history is opaque, or on transparent rails that cannot protect sensitive data. There is no environment where agents can evolve over time with an immutable audit trail while keeping their prompts, policies, and datasets private. For enterprises and regulators, this is the core blocker to deploying autonomous systems at scale.

Missing Rails for Agent-to-Agent Value Exchange

Even as autonomous agents proliferate, there is no native, neutral infrastructure for agent-to-agent value exchange and coordination. Payments, metering, and settlement between agents are bolted on rather than built in.

In areas like trading and prediction markets, most users lack access to high-quality information and systematic decision tooling. Today's platforms rarely embed agents that can continuously digest data, reason about probabilities, and act deterministically on a user's behalf, closing the information gap between retail participants and sophisticated actors.

Part 1: How Amadeus solves it

- ☑ Amadeus makes AI execution native, scalable, and verifiable by embedding agent execution directly into a high-performance Layer 1.
- ☑ Developers and organizations gain a unified environment to build, run, and evolve agents on-chain without custom infra or glue code.
- ☑ Nova AI and Agent Studio remove the need for low-level coding, while the underlying consensus and runtime guarantee deterministic behavior.
- ☑ At the same time, Agent Commerce & Privacy Rails provide built-in rails for agent-to-agent payments, metering, and settlement, so the economics of the agent economy live directly on the chain.

Part 2: Amadeus Architecture

NETWORK DESIGN

Amadeus currently has two reference nodes, one coded in Rust and another in Elixir, a choice that balances speed of code execution with speed of development. Elixir is a functional language, perfect for high throughput network code and quick software development, while Rust ensures state of the art security, performance and safety. This architecture delivers a network with a block finality targeted at under 0.5 seconds.

USEFUL PROOF OF WORK

The consensus mechanism, Useful Proof of Work, transforms mining into meaningful AI computation. Miners perform Tensorcore MatMul operations, the same tasks used in agent inference, generating real-world value while securing the network. Each epoch issues AMA according to an emissions schedule that starts with a linear bootstrapping phase and, after the Shenron Upgrade, transitions into a smooth logarithmic decay toward a 1B AMA supply cap. This preserves early network growth while progressively slowing inflation and funding a TEE-backed privacy layer. Security and scalability are reinforced by BLS12-381 cryptographic signatures. The environment is fully decentralized and focused solely on useful computation.

DETERMINISTIC EXECUTION LAYER

Amadeus offers a permissionless runtime where autonomous agents can operate with guaranteed reproducibility. On-chain agent memory and real-time coordination are built in, and execution is deterministic to ensure identical outcomes given the same inputs. Every execution path is cryptographically verifiable, creating a predictable and provable environment for intelligent applications.

Amadeus delivers speed, scale, and security through its modern stack. Its consensus model turns mining into real AI computation, and its deterministic runtime ensures agent behaviour is always verifiable. The result is a blockchain that evolves agents safely, in real time.

While the architecture delivers the foundation for speed, scale, and security, Nova AI is the intelligence layer that turns those capabilities into a seamless, verifiable agent-building experience.

TEE Pool & Private Agent Execution

As the network matures, a portion of emissions and fees will be allocated to a TEE pool: miners who run workloads inside GPU-based trusted execution environments. These TEE miners provide hardware-level isolation for sensitive agent logic and data, enabling:

- Confidential inference and decision-making
- Secure multi-party computation between agents
- Verifiable private data processing

The same consensus that powers UPoW also verifies TEE-produced outputs, so agents can execute privately while still inheriting Amadeus's deterministic behavior and cryptographic guarantees.

Part 2: Amadeus Architecture cont.

Network status (as of December 2025)

74 validators/nodes across 11PoPs

Approx 205 million transactions processed since genesis (≈867 transactions per day on average)

Current block height: 44,658,959+ blocks.

1.3 million AMA distributed per day in mining fees

16k+ agents created on-chain, with 20 verified/curated agents

12k+ connected accounts and 700k+ on-chain messages between agents and accounts

Community footprint

7.7k Discord members

7.8k Telegram members

7.4k followers on X

37k -strong Bitte.ai community now being integrated into Amadeus

Part 3: NOVA AI Architecture

SEMANTIC DEVELOPMENT ENVIRONMENT

Nova AI is a platform powered by Amadeus that allows training and running advanced AI models as part of the UPoW system. Since UPoW is supported natively in the Amadeus protocol, all nodes are incentivized to contribute their GPU resources to power up both distributed training and inference.

Nova AI continuously refines generated agents. Each new version is compiled, tested, and recorded on-chain, producing an immutable learning history that shows how and why an agent changed. Combined with TEE-backed execution, this lets organizations run self-improving agents that evolve privately, but whose evolution remains auditable and enforceable at the protocol level.

KEY PILLARS

The platform is underpinned by a versioned, immutable development graph. All logic artifacts are stored as semantic nodes, and all modifications are tracked. The Context Assembler constructs rich prompts for LLM pipelines, making the process of refining and compiling agents more efficient. By integrating test-first methodology, Nova guarantees correctness by design and builds a foundation for agents to be both safe and adaptable.

Nova AI transforms how software is built. Developers describe the desired outcome, and Nova orchestrates the rest, from planning and testing to deployment. By combining semantic structure with automated refinement, it delivers agents that are not only functional but also traceable and verifiable from inception.

Part 4: Tokenomics

Token Overview

The AMA token has a projected total supply of around 1 billion over an estimated 15 to 30 years. Based on current emissions, supply is expected to approach 1 billion AMA at around year 20. As of December 2025, the circulating supply stands at roughly 450 million AMA. The launch model distributes 100% of the supply through Useful Proof of Work (UPoW), with no VC allocations and no premine. This ensures that the tokenomics are fully aligned with those who build, mine, and contribute to the network.

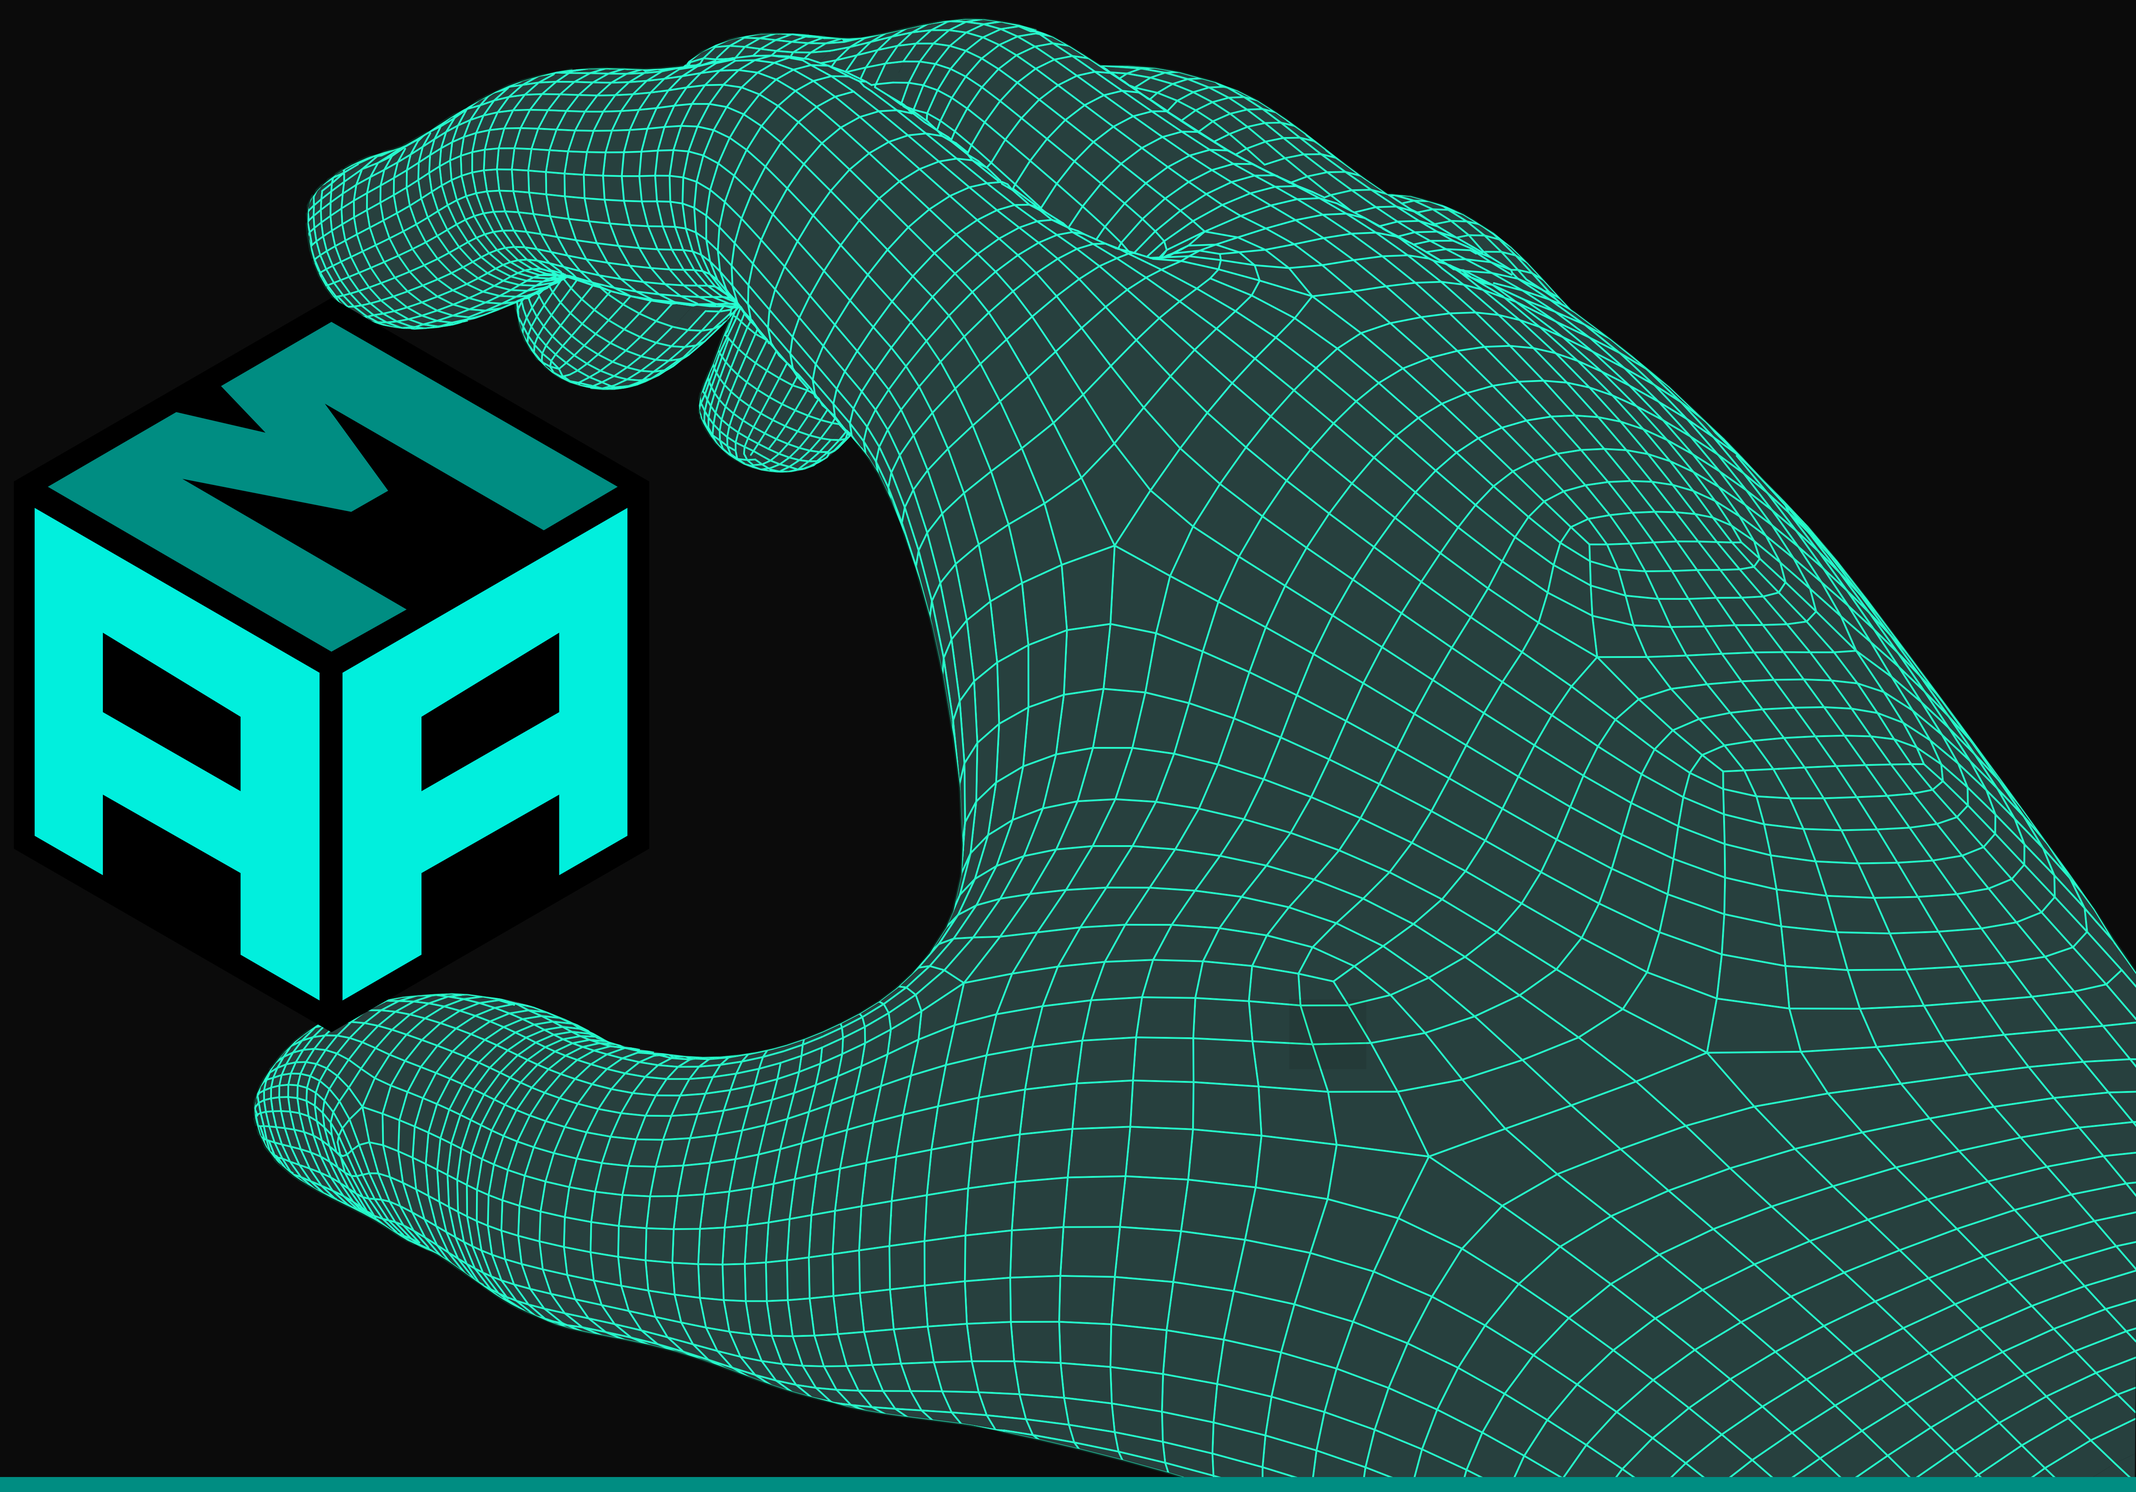

Part 4: Tokenomics cont.

Emissions Schedule and Halving

Amadeus uses a three-phase emission model designed to bootstrap security, align miners with long-term sustainability, and fund privacy infrastructure for agents:

From epoch one there was a curve issuing on average 1M tokens per epoch and gradually reducing to 868k AMA per epoch at epoch 422 and then to 779k AMA by epoch 749. At epoch 750, the Shenron Upgrade (Halving Event), the network transitions from the current curve to a smooth logarithmic decay function, marking a significant shift in tokenomics. This transition reduces per-epoch issuance by approximately 65% (from 779k to 271k AMA), slowing token release while maintaining predictable supply growth toward the 1 billion AMA target. The logarithmic curve prevents abrupt inflation shocks and supports long-term sustainability, with issuance continuing to taper gradually after epoch 810 until total supply of 1B AMA is reached.

Concurrent with the curve transition, the mining fee begins decreasing at epoch 750, dropping from 85% to 25% over 60 epochs (1% per epoch reduction). This creates a powerful incentive realignment: while total issuance decreases, miner rewards actually increase from 43k AMA at epoch 750 to 186k AMA by epoch 810, exceeding previous miner earnings despite lower overall emissions. The additional miner allocation during epochs 750-810 will be earmarked for the upcoming TEE pool split, where miners operating Trusted Execution Environment enclaves will receive premium rewards for contributing privacy-preserving compute infrastructure. This dual mechanism ensures robust network security while accelerating the buildout of infrastructure for the agent-to-agent economy.

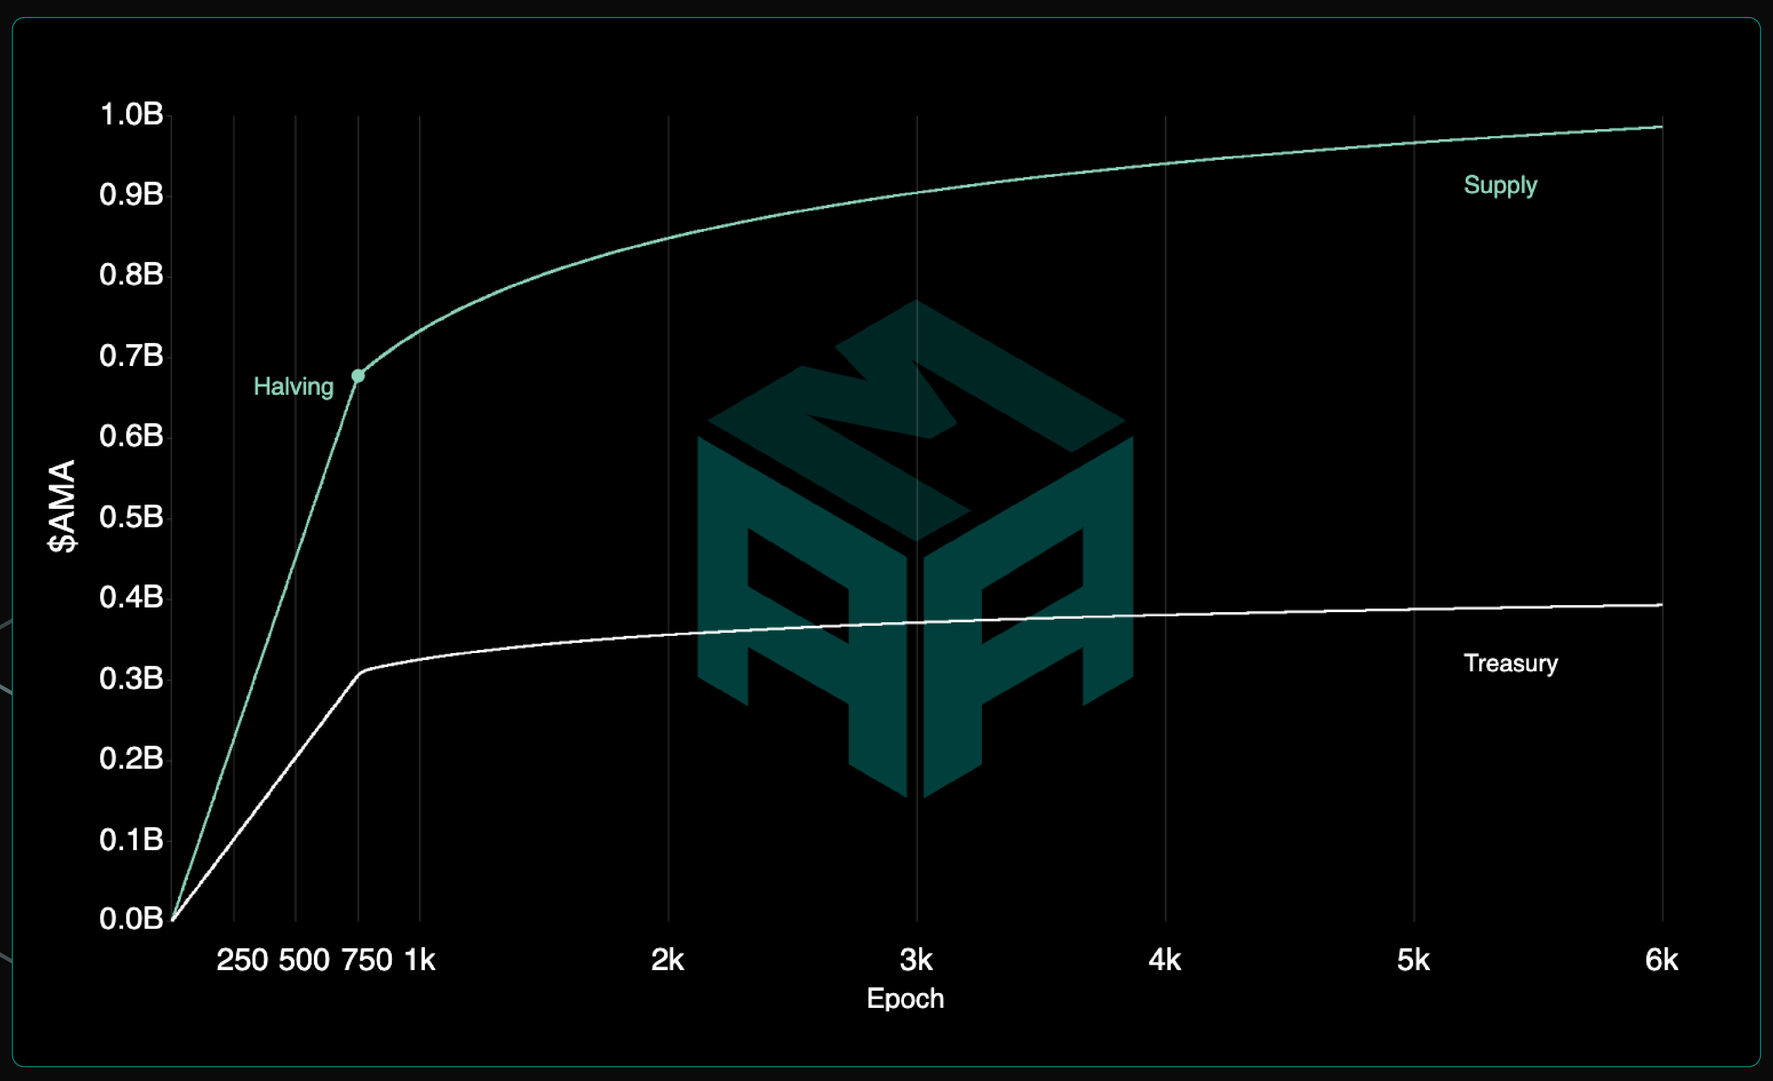

Part 4: Tokenomics cont.

Economic Rationale

The emission schedule of $AMA follows a gradual decay model designed to balance fairness, sustainability, and long-term network growth. Rather than concentrating supply in early adopters, the model ensures that rewards decrease progressively as adoption and utility increase. This creates a natural transition from incentivized bootstrapping to organic demand-driven value.

Early participants are rewarded for taking higher technical and market risk, while later participants still benefit from predictable token availability for staking mechanisms, AI agent deployment, and governance. The smooth decay also reduces inflationary pressure over time, supporting a more stable market environment and aligning incentives between miners, builders, and token holders.

A portion of emissions and fees in the mature phase is earmarked for TEE-backed privacy infrastructure, ensuring that as the agent economy scales, there is a sustainable budget for confidential, high-assurance compute.

In essence, the model prevents supply shocks, encourages continuous participation, and ensures that no single period of the network's history captures disproportionate value, a foundation for long-term decentralization and equitable distribution.

Part 4: Tokenomics cont.

Token Utility

AMA serves multiple purposes within the ecosystem: it fuels AI agent execution, incentivizes UPoW miners and covers fees for transactions and network operations. It also underpins future governance participation.

The AMA token serves exclusively as a functional utility instrument within the Amadeus network ecosystem, enabling essential network operations. AMA is not a security, investment contract, equity interest, or profit-sharing arrangement, and holders should have no expectation of profit or financial return. AMA facilitates three core network functions:

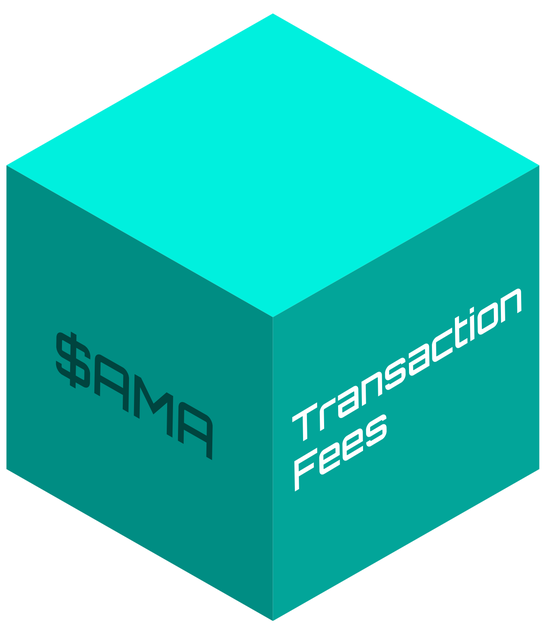

Network participants expend AMA to pay fees for asset transfers, smart contract deployments, and on-chain operations. These fees compensate validators and miners who provide computational resources and secure the blockchain infrastructure, functioning similarly to gas fees on other blockchain networks.

AMA tokens are consumed to power AI agent operations on the network, including computational resources for agent inference, on-chain memory storage, and state transitions. Users expend AMA to deploy and run autonomous agents, with costs denominated based on computational requirements and resource consumption necessary for network operations.

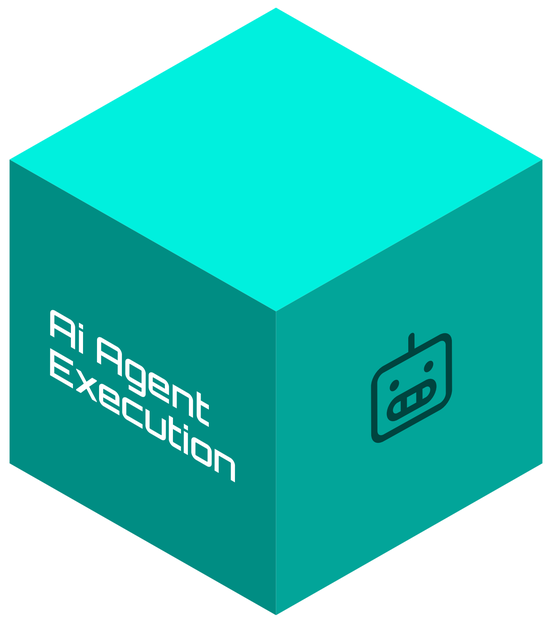

AMA is distributed as compensation to UPoW miners who provide computational resources by performing Tensorcore MatMul operations that secure network consensus. Mining rewards compensate participants for hardware costs, electricity expenditure, and infrastructure provision necessary for network operation, similar to mining rewards in other proof-of-work blockchain systems.

Part 4: Tokenomics cont.

Amadeus Agent use cases

Amadeus enables deployment of autonomous AI agents that operate directly on-chain, coordinate with each other, and tap into external services while inheriting the platform's verifiable execution layer. With Agent Commerce & Privacy Rails and Nova AI, these agents are first-class economic actors in the emerging agent-to-agent economy. Examples include:

Agent Type	What the Agent Does	Notes
DeFi Copilot Agents	Automate portfolio actions, monitoring and execution logic for DeFi workflows.	Runs instructions generated by Nova; executes actions through Amadeus.
x402 Commerce Agents	Access APIs, pay for services, interact with external systems using HTTP 402 payment semantics.	Use x402 flows; AMA may be used if designated as a payment asset.
Compliance & Monitoring Agents	Watch wallets, flows, protocols, sanctions lists; produce automated compliance outputs.	Generate on-chain or off-chain logs; Nova builds logic.
Prediction Market Strategy & Settlement Agent	Act as embedded "copilots" and strategy managers for prediction markets (e.g. PredX-style platforms). They analyze market questions, odds, liquidity, and external data feeds, then generate forecasts, suggested bet sizes, or fully automated strategies across multiple markets.	Deterministic logic makes every trading or settlement decision replayable and auditable. Agents can pay for data, news, and model inferences over x402-style HTTP payments, and keep proprietary alpha and risk models private inside TEEs, while exposing only verifiable on-chain actions.
Private Strategy Agents	Runs proprietary decision logic (trading, pricing, bidding, or policy enforcement) on private data inside TEEs	Executes in TEE-backed privacy pools; evolution tracked by Nova; only outputs are revealed on-chain.
Robotics Fleet Orchestration Agent	Coordinates fleets of robots (factory arms, warehouse bots, delivery drones) by assigning tasks, routing paths, and enforcing safety policies. The agent ingests sensor data, job queues, and environment state, then deterministically decides which robot should do what, when, and where.	Runs as a deterministic Amadeus agent so decisions are reproducible and auditable for safety. Can learn from throughput and incident data via Nova AI, while keeping raw sensor and operational data private inside TEE-backed execution environments.

Flagship Vertical: "Stellaris" Private Agentic Trading

The core of Stellaris is formed together by the Prediction Market Strategy & Settlement Agents and Private Strategy Agents. In Stellaris, TEE-secured agents consume on-chain and off-chain market data, execute deterministic strategies across CEX, DEX, DeFi, and prediction markets, and settle flows over x402-style payment rails, while keeping proprietary models and risk logic confidential.

Part 5: Consensus – Useful Proof of Work + BFT

Amadeus combines Useful Proof of Work (uPoW) with a Byzantine Fault Tolerant (BFT) consensus mechanism to maintain a fast, resilient, and trust-minimized ledger. Each block, called an entry, represents roughly 0.5 seconds of real-world time. Consensus blends real computation with cryptographic voting, achieving sub-second finality while performing useful AI tasks.

Part 5: Consensus – Useful Proof of Work + BFT cont.

Core components

Term	Definition
Peer	Any node running the Amadeus protocol and participating in the network.
Validator	A validator is selected through UPoW score and performance. Validators propose blocks and collect network fees.
Entry	A block containing transactions, representing ~0.5 seconds of real-world time.
Epoch	A sub-ledger of 100,000 entries (~14 hours), during which a fixed set of trainers operate.
Attestation	A trainer's cryptographic vote on an entry's validity.
Aggsig	An aggregate BLS12-381 signature from > ⅔ of trainers, proving finality for an entry.
Contractstate	The shared state that includes balances, smart contracts, and the active trainer set.
Catchup	The process that a node uses to fetch missed entries, attestations, and signatures after downtime.

Part 5: Consensus – Useful Proof of Work + BFT cont.

Consensus Flow

Each round follows three tightly coupled phases:

1 Proposal & Gossip

A designated trainer bundles valid transactions and gossips the proposed entry to the network.

↓

2 Verification & Voting

All nodes verify the block's useful-work proofs—matrix multiplication results that double as AI computation—and broadcast cryptographically signed attestations.

↓

3 Aggregation & Finality

BLS12-381 signatures are aggregated into a single proof. Once a super-majority (> 67 %) is reached, finality is achieved in ~0.5 seconds, and the entry is irrevocably committed.

Part 5: Consensus – Useful Proof of Work + BFT cont.

Work → Verification → Mining

Beyond consensus mechanics, Amadeus integrates real-world computation directly into block production:

1 Work

Nodes solve an AI task (matrix multiplication) and submit it on-chain.

2 Verification

Other trainers verify the solution and award a point to the solver.

3 Mining

Trainers with the highest points earn the right to propose blocks and collect transaction fees by becoming validators.

Over time, a subset of these AI tasks will be routed into TEE enclaves, turning the same mining process into a source of privacy-preserving compute for agents, with premium rewards for GPU-TEE miners.

Trainer Slot Formula

At the start of each epoch, trainers are assigned slots that determine how many entries they must propose. If a trainer fails to propose in their slot, they risk slashing.

$$\text{entry\_epoch} = \frac{\text{entry\_height}}{100{,}000}$$

$$\text{entry\_slot} = \text{entry\_height}$$

$$\text{trainer}_i \quad \text{where} \quad i = \text{entry\_slot} \bmod \text{epoch\_trainers}$$

Where i is the trainer's index in the epoch trainer list.

Part 5: Consensus – Useful Proof of Work + BFT cont.

Aggregate Signatures

Consensus finality relies on BLS aggregation. If at least 67 % of trainers sign the same entry, the aggregate signature proves consensus:

$$\text{aggsig} = \sum_{i=0}^{\lfloor 0.67\,\text{epoch\_trainers} \rfloor} \text{bls}(\text{entry},\ \text{trainer}_\text{i})$$

This allows the network to tolerate up to 33 % faulty trainers while maintaining safety and liveness.

uPoW Difficulty Adjustment

Each epoch dynamically adjusts solution difficulty to match network UPoW speed. A valid solution's hash must begin with a certain number of zero bits determined by the previous epoch's performance:

$$\text{next\_bits} = \begin{cases} \max\big(20,\ \min(\text{prev\_bits} - 3,\ 64)\big), & \text{if sols} = 0 \\ \max\big(20,\ \min(\text{prev\_bits} + \delta,\ 64)\big), & \text{if sols} > 418{,}000 \\ \max\big(20,\ \min(\text{prev\_bits} - \delta,\ 64)\big), & \text{if sols} < 342{,}000 \\ \text{prev\_bits}, & \text{if } 342{,}000 \leq \text{sols} \leq 418{,}000 \end{cases}$$

Part 5: Consensus – Useful Proof of Work + BFT cont.

Where the factor of difficulty adjustment is:

$$\delta = \begin{cases} \min\Big(\max\Big(\lceil \frac{1}{2} \lceil \log_2(\text{sols}/380{,}000) \rceil \rceil,\ 1\Big),\ 2\Big), & \text{if sols} > 418{,}000 \\ \\ \min\Big(\max\Big(\lceil \log_2(380{,}000/\text{sols}) \rceil,\ 1\Big),\ 3\Big), & \text{if sols} < 342{,}000 \end{cases}$$

Slashing & Special Business Trainer

Validators rotate every epoch, but the validator can also be slashed mid-epoch by one of the validators manually triggering the slashing mechanism for:

- Proposing two entries at the same height
- Delaying entry creation (> 8 s mid-epoch / > 30 s end)
- Taking > 0.6 s on average to create entries
- Being offline and not proposing an entry

Part 5: Consensus – Useful Proof of Work + BFT cont.

Emissions & Rewards

Until epoch 750, each epoch issues emissions distributed as follows:

$$\text{epoch\_emission} = \frac{11{,}536{,}480{,}000}{(\text{epoch} + 690.573766)^{3/2}}$$

$$\text{p67\_trainer\_reward} = \left(\frac{1}{67}\right) \times \frac{6 \times \text{epoch\_emission}}{7}$$

At epoch 750, Amadeus transitions to a different function, aka. Shenron Upgrade.

$$\text{epoch\_emission} = \frac{1{,}430{,}936{,}428}{(\text{epoch} - 20)^{1,3}}$$

Part 5: Consensus – Useful Proof of Work + BFT cont.

Catchup

Nodes that fall behind can synchronize through catchup, downloading entries, attestations, and aggregate signatures in fixed-size chunks. If too far behind, they reset from a merkle-state-rooted snapshot. This ensures resilience in real-world conditions.

Summary

Amadeus consensus integrates real computation, BFT finality, and dynamic trainer selection into one mechanism. With sub-second finality, adaptive difficulty, and mathematically provable safety, it creates a network that's both fast and meaningfully productive.

** The current consensus mechanism is live, but useful mining (full UPoW workloads routed to AI/TEE tasks) is being phased in (see roadmap).

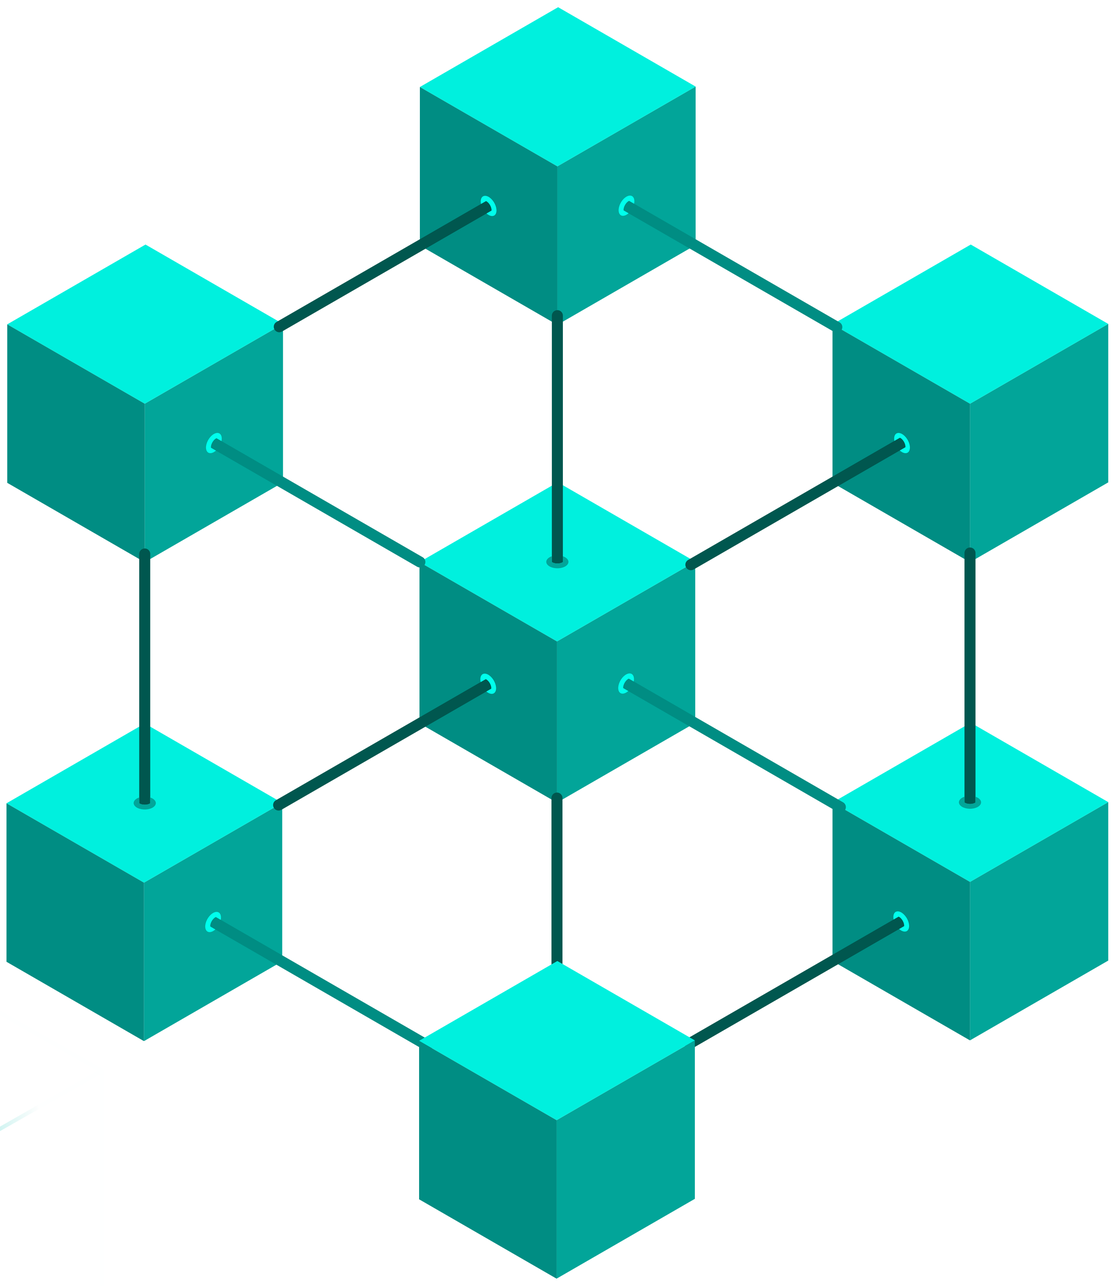

Part 6: Roadmap

Q2 2025 — ☑ Delivered

Public Testnet, Mainnet Launch, WASM Smart Contracts, and Blockchain Infrastructure Deployment — the foundations for Amadeus have been delivered, enabling developers to begin building with the full agent runtime from day one.

Q3 2025 — ☑ Delivered

Website Launch, Self-Custody Wallet, and Block Explorer, providing the essential user-facing tools for accessing and interacting with the network.

Q4 2025 — ☑ Delivered

Nova Inference POC, protocol hardening, Genesis hackathon.

Q1 2026

Litepaper release; Stellaris Agent Studio + Launchpad; Useful Proof of Work (uPoW) mining; ecosystem grants; Prime Points rewards program; browser extension + mobile wallet.

Q2 2026

Nova Network (agent orchestration); x402 payment rails; TEE privacy layer; yield-locking vaults (optional / launch partners).

Q3 2026

Enterprise Solutions Launch and Agent Commerce Protocol, targeting enterprise adoption and funding ecosystem growth.

Q4 2026

GPU marketplace; agent economy primitives; Nova privacy upgrades; Stellaris scale upgrade, Governance v1.

The roadmap moves from infrastructure to adoption, ensuring each phase builds on the last to scale compute capacity, liquidity, and practical use of autonomous agents.

Part 7: Market Opportunity

Market Growth

The AI Agents market is projected to grow from $7.92 billion in 2025 to $236 billion by 2034 (45.82% CAGR)[1], while the Agentic AI market is expected to expand from $13.81 billion in 2025 to $140.8 billion by 2032 (39.3% CAGR)[2]. These figures highlight the rapid acceleration of autonomous system adoption across industries.

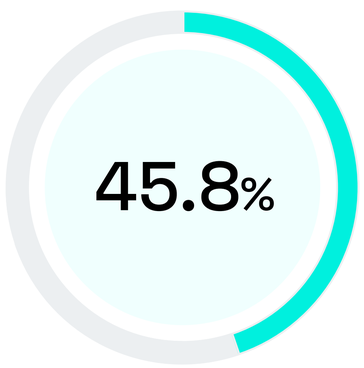

AI Agents Market CAGR

$7.92B → $236B

(2025–2034)[1]

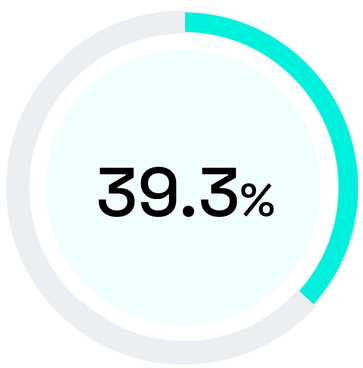

AI Infrastructure CAGR

$13.81B → $140.8B

(2025–2032)[2]

Competitive Analysis

The convergence of AI and blockchain networks is fueling a new generation of decentralized networks that power inference, model training, and agent execution. Amadeus sits at the intersection of AI compute networks and agent execution. Emerging category leaders typically focus on a single layer of the stack:

Bittensor (TAO)

Incentivizes model training and inference through its subnet marketplace; rewards models and subnets rather than end-to-end agents.

Gensyn

A decentralized marketplace for machine-learning training jobs, using blockchain verification for training workloads, without a native agent runtime.

io.net

A decentralized GPU cloud for AI workloads and rendering, focused on GPU rental and compute payments rather than persistent agents.

OG

An AI-first modular Layer 1 that combines scalable execution with a compute network, but is not opinionated about agent lifecycle or consensus-verified agents.

SingularityNET

A cross-chain AI services and model marketplace that coordinates APIs and models, not consensus-level agent execution.

Sources

Precedence Research — AI Agents Market Size, Share, Growth (2024)
https://www.precedenceresearch.com/ai-agents-market

MarketsandMarkets — Agentic AI Market Worth USD 140.80 Billion by 2032
https://www.prnewswire.com/news-releases/agentic-ai-market-worth-140-80-billion-by-2032-exclusive-report-by-marketsandmarkets-302490696.html

Part 7: Market Opportunity cont.

In parallel, a set of app-layer / non-L1 projects focus on agent runtimes and UX rather than base-layer consensus:

Virtuals

A society of productive AI agents with a strong SDK and tokenization platform, but app-level monetization only and no consensus-verified agent layer.

Eliza (elizaOS)

An open-source framework and “agent OS” for building and orchestrating agents, typically running off-chain or on hosted infrastructure.

GaiaNet / Gaia

A decentralized agent hosting network that lets developers deploy and monetize AI agents across nodes, with state local to each node rather than globally verified.

As summarized in the comparison table above, most of these networks optimize for raw compute (training or inference), SDKs, or app-layer UX, but they do not yet deliver the full spectrum of native agent runtime, persistent protocol-level agent state, consensus-verified agent execution, and agent-centric incentives in a single Layer 1. Amadeus is the first blockchain where mining directly contributes to agent work, where agents have protocol-level persistent state, and where developers are rewarded based on verified agent usage, with Nova AI providing a full deterministic lifecycle for self-improving agents.

Amadeus is the first blockchain where mining contributes to AI evolution, and where anyone can build, deploy, and scale self-improving AI agents natively without code. The on-chain agent evolution is enabled by the Nova AI compiler which allows agents to self improve with immutable on-chain learning history, powered by inference from UPoW consensus computations.

By unifying compute, intelligence and execution within a single layer 1, Amadeus transforms decentralized infrastructure into the world’s first thinking blockchain with evolving intelligent agents, laying the foundation for the next era of autonomous AI systems.

Part 7: Market Opportunity cont.

Go-to-market Strategy

Amadeus employs a phased community-first approach designed to build authentic engagement before scaling to mass adoption. The foundation phase focuses on establishing core infrastructure and creating high-quality educational content that positions Amadeus as the ”first thinking blockchain”, emphasizing both instant asset transfers and no-code AI agent deployment. This narrative is amplified through strategic micro and mid-tier influencer partnerships within AI and crypto communities, supported by targeted advertising across Telegram, Twitter and Reddit.

Target Builders: Utility & Infrastructure Agents First

Within the broader agentic economy, Amadeus focuses primarily on utility and infrastructure agents, agents that touch real capital, data, and reputation.

We distinguish between two broad segments:

Social / culture agents

Agents anchored in lore, personalities and branding. These are primarily driven by narrative and community, and often experiment with token-based speculation.

Utility / infrastructure agents

Agents that execute concrete workflows such as DeFi strategy execution, treasury and risk co-pilots, compliance and monitoring agents, and deterministic prediction market agents. These agents are judged on credibility, risk management, reliability, and execution quality.

Amadeus can support both segments, but our go-to-mark is focussed on utility and infrastructure agents that:

- **require deterministic and auditable execution,**
- **benefit from privacy-preserving computation over sensitive data and proprietary logic, and**
- **need native machine-to-machine payments to coordinate value flows at scale.**

This focus aligns directly with Amadeus’ deterministic runtime, TEE-backed privacy layer, and x402-style commerce rails, positioning the network as the home for agents that behave as **critical infrastructure** rather than purely social assets.

Part 7: Market Opportunity cont.

Who are Amadeus' ideal builders?

Amadeus is built for teams that need deterministic, self-improving agents to touch real money, real users, or real infrastructure. In practice, our "perfect customers" are:

PredX-style prediction markets and protocols that want AI copilots, strategy vaults, and AI "house" agents whose behavior is replayable and auditable.

Dexes and onchain exchanges that need risk engines, liquidity managers, and MEV-aware routing agents which must behave deterministically and can evolve safely over time.

CEXs and brokerage platforms that want to offload market-making, surveillance, and routing logic to verifiable agents while keeping strategies and data private inside TEEs.

Yield optimizers, structured products, vaults, and DAOs that need agents to manage portfolios, collateral, and risk with a full on-chain learning history.

Part 7: Market Opportunity cont.

Flagship Vertical: “Stellaris” Private Agentic Trading

To make these ideas concrete, Amadeus is launching Stellaris, a flagship vertical for private agentic trading built on top of the core protocol.

Stellaris targets the same builders described above: prediction markets, DEXes and onchain exchanges, CEXs and brokerages, yield and treasury managers, but gives them a shared, deterministic and privacy-preserving execution environment:

Unified market state

Agents consume a CRDT-style, unified view of state from CEXs, DEXs, DeFi protocols, and prediction markets instead of stitching together fragmented feeds themselves.

TEE-backed privacy

Trading and risk logic runs inside TEEs, so proprietary alpha and risk models remain confidential while actions and outcomes stay replayable and auditable on-chain.

Deterministic, verifiable behavior

Every Stellaris agent runs on the Amadeus deterministic runtime, making each decision reproducible from inputs and versioned agent logic, closing the information gap between sophisticated actors and everyone else.

Stellaris is Amadeus’ initial go-to-market wedge: it demonstrates how private, deterministic agents can manage real capital in production and serves as a template for future verticals in DeFi, infrastructure, and enterprise workflows.

STELLARIS
PRIVATE AGENTIC TRADING

Part 7: Market Opportunity cont.

Traction to date

Amadeus mainnet has been live since March 2025. As of December 2025, the network has processed over 205 million transactions (≈867 per day) across 44.6 million+ blocks, with 16k+ agents, 12k+ connected accounts, and 700k+ on-chain messages. The protocol is currently secured by 74 validators/nodes across 11PoPs. On the ecosystem side, the combined community spans roughly 7.7k Discord, 7.8k Telegram, and 7.4k X followers, plus a 37k-strong Bitte.ai community base, following Amadeus' 2025 acquisition of the Bitte.ai agent infrastructure product, whose technology and users are being folded into Nova AI and Agent Studio.

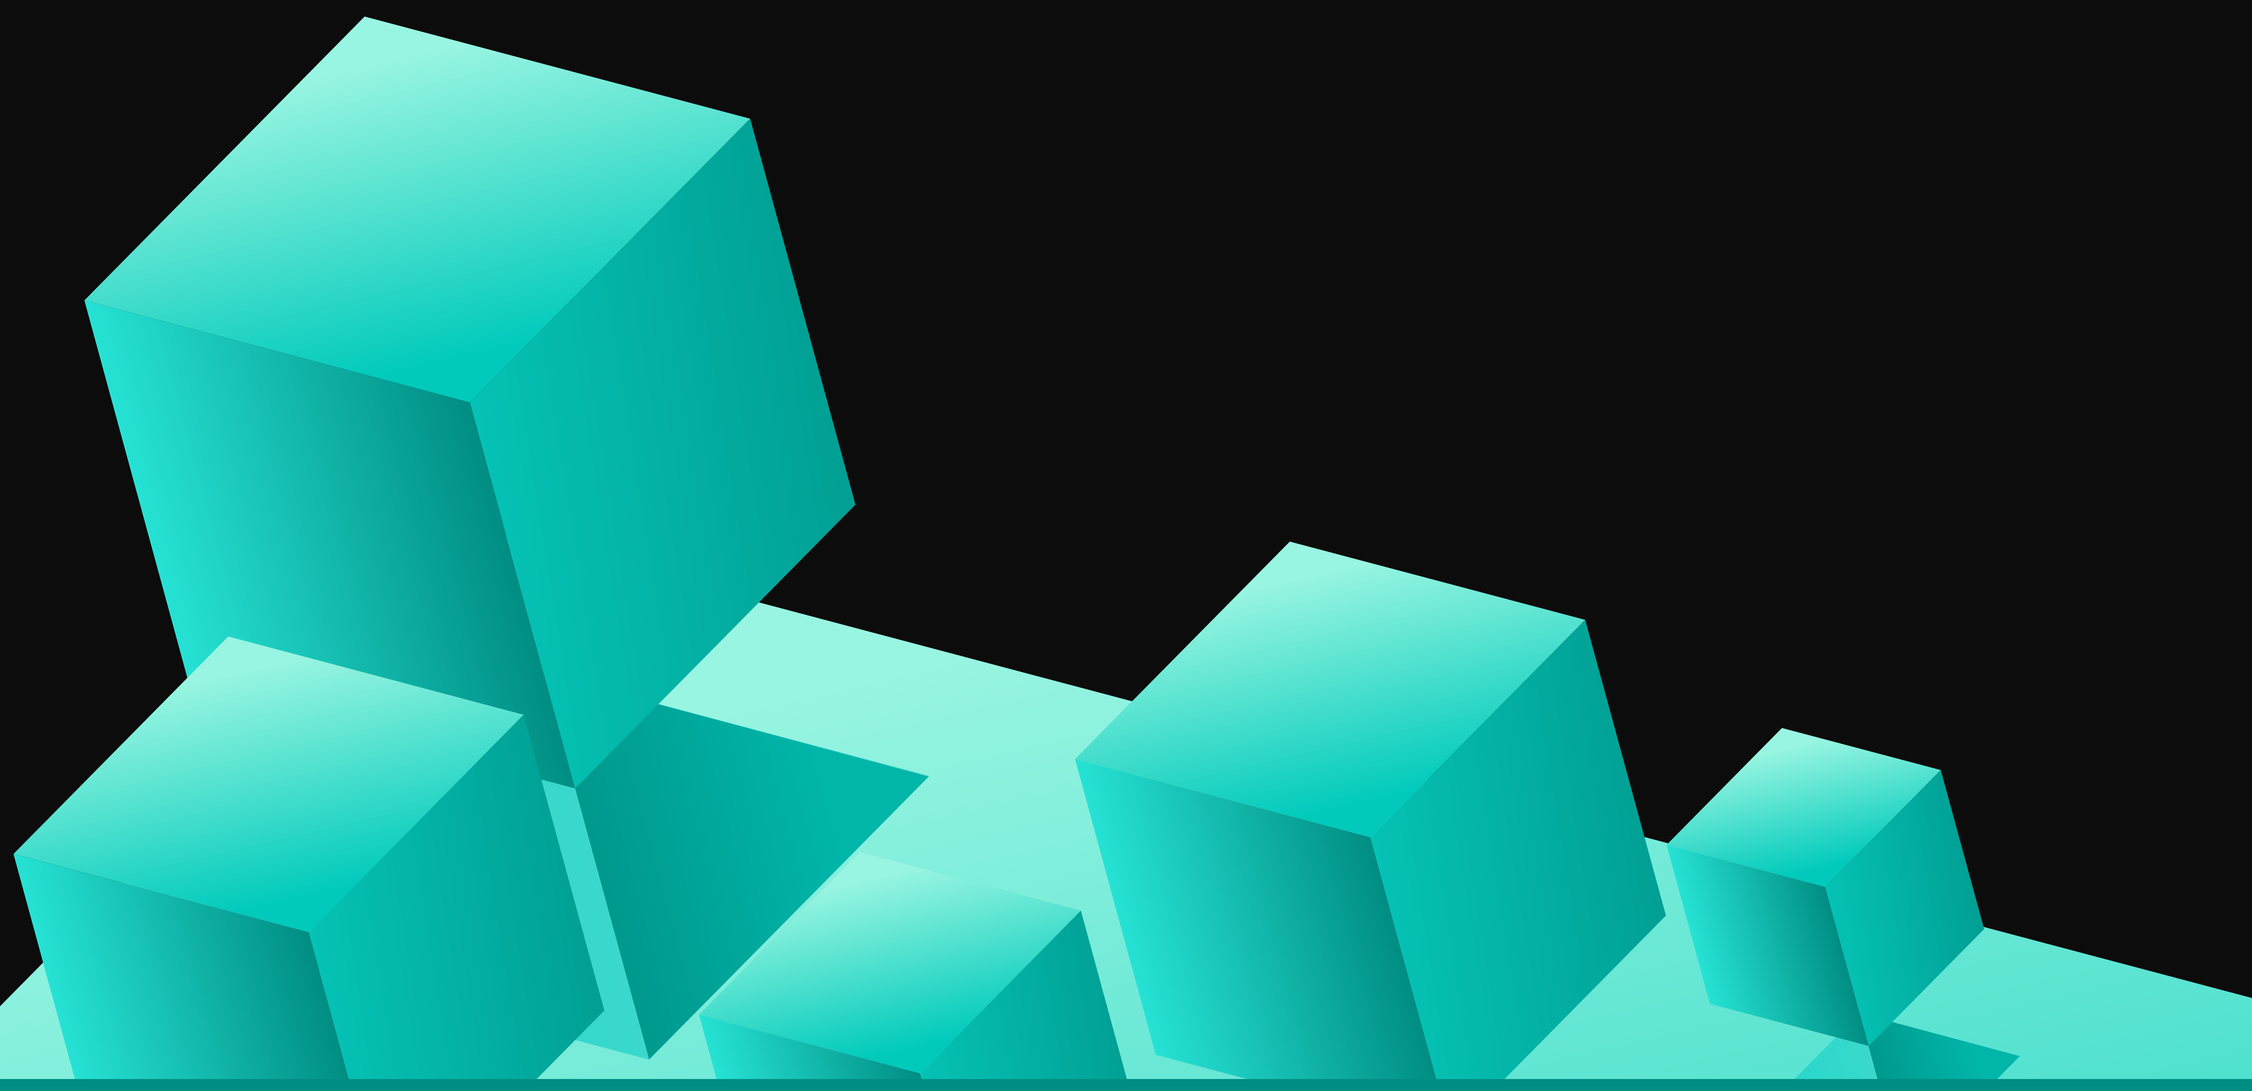

Part 7: Market Opportunity cont.

Default Agent Monetisation Model

A core element of the go-to-market strategy is how agents earn and share value on Amadeus.

By default, agents monetize through fees and revenue in AMA and stable assets, using:

- **usage-based fees for calls and workflows,**
- **performance and success fees for financial strategies, and**
- **subscription or pay-per-call models routed through x402-style payment flows.**

Revenue can be programmatically shared between:

- **the agent's creators and maintainers,**
- **integrators and distribution partners, and**
- **other contributors, via on-chain rev-share contracts.**

This default model allows agents to be economically viable from day one without requiring their own fungible token, while keeping user experience and liquidity concentrated around a small set of core assets, with AMA as the primary economic engine.

Position on Agent Tokens

On Amadeus, agent-specific tokens are optional, not the default.

For the majority of utility and infrastructure agents, introducing a separate tradable token would fragment liquidity, complicate user experience, and risk misaligning incentives toward speculation rather than robust, risk-aware performance. Instead, most agents are intended to operate as services that charge fees and settle revenue flows in AMA and other liquid assets.

Agent-specific tokens are reserved for advanced cases where an agent effectively behaves like an independent protocol, for example:

- **complex, multi-stakeholder governance,**
- **explicit risk-backstopping via staking and slashing, or**
- **long-term, independent treasury management.**

Even in these cases, AMA remains the base asset for gas, settlement, and cross-agent payments, ensuring that value continues to accrue to the broader Amadeus economy while allowing specialized agents to express their own governance and risk surfaces where appropriate.

Part 7: Market Opportunity cont.

Design Partner and Ecosystem-Led Motion

In the first phase, Amadeus prioritizes design partners across L1/L2 ecosystems, DeFi protocols, and infrastructure providers. These partners co-design high-value agents such as risk bots, embedded market copilots for PredX style dApps, liquidity managers, and privacy-preserving monitoring agents using Nova AI and Agent Studio. This creates concrete, referenceable case studies and production workloads that validate the network's capabilities and inform product development. It's closing the information gap by digesting data, pricing probabilities, and executing deterministic strategies on behalf of users, while paying for external data and signals through x402-style HTTP payments.

In parallel, Amadeus launches a developer and ecosystem program that provides grants, technical support, and templates for teams building agents on the network. Agent Studio and the underlying Agent Commerce & Privacy Rails are designed to make it easy for builders to plug into the agent-to-agent economy without building custom payment, metering, or privacy infrastructure from scratch.

Phased Expansion & Scale

As core use cases and design partners mature, Amadeus scales through three reinforcing motions:

3. Scale: evolve into a broadly adopted agent-to-agent economy, where wallets, protocols, enterprises, and independent developers deploy agents that coordinate, trade, and settle across Amadeus. At this stage, the focus shifts to higher-order network effects: agent marketplaces, reputation systems, privacy modules, and enterprise integrations that leverage Amadeus as the default rails for agent-native commerce.

2. Expansion: grow the number of production agents and protocols integrated with Amadeus; support x402-style HTTP-native payment flows for agents interacting with external APIs and services; and expand the ecosystem through targeted grants and hackathons with an emphasis on DeFi, prediction markets and infrastructure.

1. Foundation: solidify the protocol, Nova AI, and Agent Studio; publish open specifications for agent execution, commerce, and privacy; and seed early on-chain agents across DeFi and infrastructure.

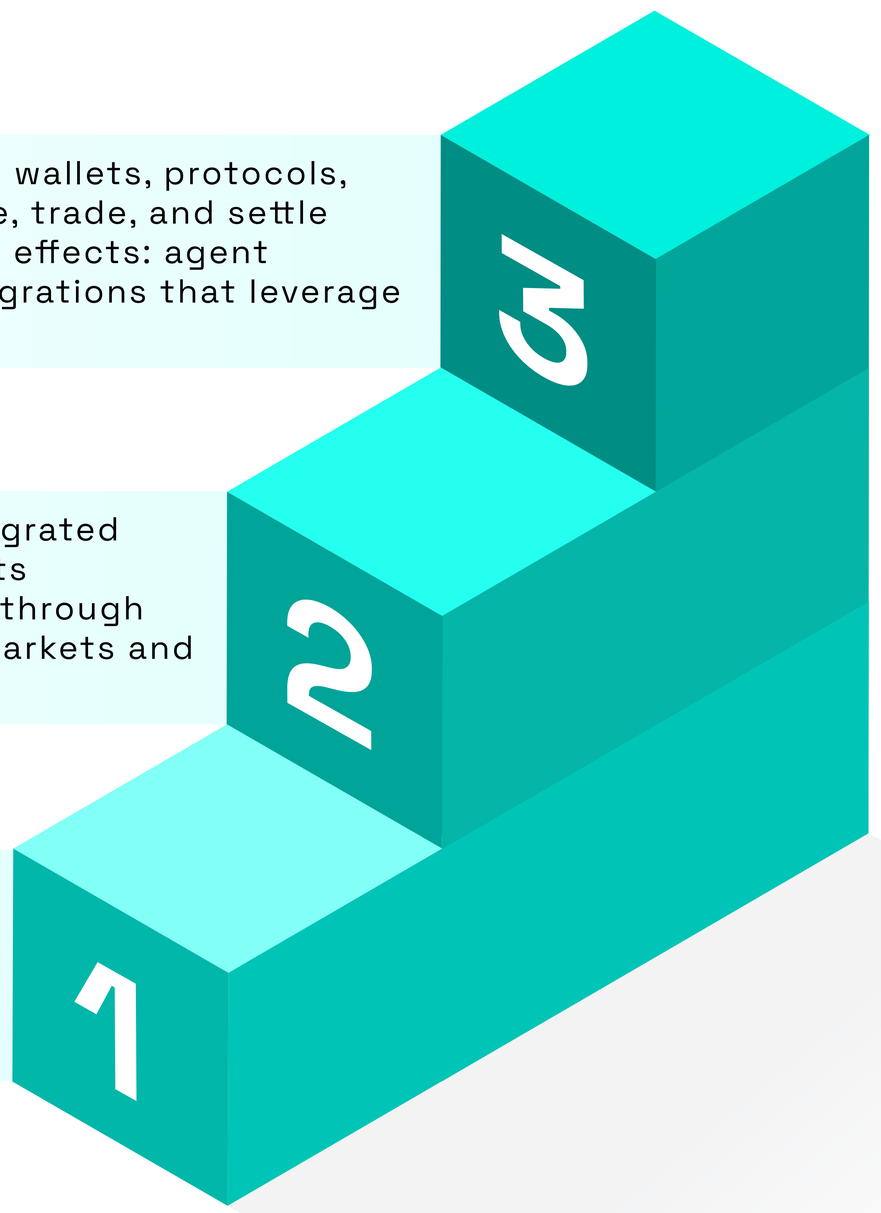

Part 8: Core Team

Ivan Svirid

Co-Founder | in

13+ years building high-concurrency systems in Erlang/Elixir and Rust; deep distributed systems and open-source background; ex NEAR ecosystem

Leslie Kivit

Co-Founder | in

15+ years building high-growth startups in and outside crypto; ex Qubic, Xapo Bank, Meta, Sennder, Booking.com, Rocket Internet

Valentyn Faychuk

Lead Backend Engineer | in

9+ years building scalable decentralized systems (Cisco, Crown Sterling, DeTEE, Sofenty)

Laura Beniguel

Lead Ecosystem | in

Ecosystem growth at major L1s (Aptos, Sei, Tezos)

Tatiana Hernandez

UI/UX Design | in

8 years experience in UI/UX and product design (Ex Moralis)

Dr. Kody Law

Advisor | in

AI Researcher & Advisor (ex Meta)

Dr. Andrea Cerone

Advisor | in

Blockchain Researcher & Engineering Advisor (Staff Researcher at Plasma)

Aaron Nevin

Advisor | in

Strategy & Growth (ex Avalanche Head of Ecosystem)

Karin Lorez

Legal Counsel | in

Founder Lorez Legal, clients: Maerki Baumann, Qubic and Zuitzerland among others

Backed by a broader team of experienced crypto developers, marketers, and ecosystem builders who have launched and managed tier 1 crypto projects. The full Amadeus team can be found on our website ama.one.

The Amadeus team blends deep protocol experience with real-world product launches. From building mining clients and smart contracts to scaling ecosystems and growth, each member brings domain-level experience in AI, blockchain, and systems engineering. This is a team that builds from scratch, and scales.

Part 9: The Vision

Amadeus is building a world where AI is not just a service but an autonomous economy - verifiable, uncensorable, and owned by those who contribute. By combining cutting-edge consensus, a deterministic execution layer, and Nova's no-code agent creation, Amadeus removes the friction between idea and intelligent execution. The network is designed to grow with its agents, compounding innovation in a transparent, open environment. This is where the next wave of AI will run - on-chain, at scale, for all. Amadeus aims to host 10K verifiable agents and 100K daily active AI transactions by 2027.

By combining deterministic execution, self-improving agents with on-chain learning histories, and privacy-preserving TEE infrastructure, Amadeus aims to become the default home for AI systems that must be both powerful and trustworthy.

Amadeus is the protocol and infrastructure to enable AI on-chain. By combining cutting-edge consensus, a deterministic execution layer, privacy TEEs, and NOVA no-code agent creation; it removes the friction between idea and intelligent execution.

The network is designed to grow and evolve the components deployed on it, whether that be an AI agent, robotic android or even frontier LLM.

Amadeus is where the next wave of AI will run - private, decentralized and open to all.

Disclaimer

LEGAL DISCLAIMER

READ CAREFULLY. CONSULT YOUR LEGAL, FINANCIAL, AND TAX ADVISORS BEFORE PROCEEDING. This litepaper is provided for informational purposes only and does not constitute an offer, solicitation, investment advice, or any contractual obligation. The Amadeus team reserves the right to modify, discontinue, or abandon the project at any time without notice or liability. AMA is a utility token designed exclusively for network operations within the Amadeus ecosystem and is not a security, investment contract, equity interest, debt instrument, or profit-sharing arrangement under any jurisdiction's laws. Purchasers should have no expectation of profit from the efforts of others or any return on investment. No representations, warranties, or guarantees are made regarding AMA token value, liquidity, market adoption, network functionality, security, development timelines, or roadmap completion. AMA tokens may have no value whatsoever and participation may result in complete financial loss. All forward-looking statements regarding development plans, market opportunities, and projections are subject to change or cancellation without notice, and actual results may differ materially. Cryptocurrency and blockchain regulations vary significantly by jurisdiction and are subject to rapid change. Participants are solely responsible for determining and maintaining legal compliance in their respective jurisdictions. AMA tokens are not available to U.S. persons (including citizens, residents, and green card holders) or residents of any jurisdiction where cryptocurrency transactions are restricted or prohibited. The Amadeus team reserves the right to implement KYC/AML procedures and refuse service to any person or entity at its sole discretion. Blockchain technology involves substantial inherent risks including but not limited to software vulnerabilities, cyberattacks, consensus failures, mining attacks, smart contract exploits, permanent loss or theft of private keys, network forks, quantum computing threats, and adverse regulatory developments. Users assume all technical, security, operational, and financial risks associated with the Amadeus network and AMA tokens. Market data and third-party information contained herein have not been independently verified and should not be relied upon. Participants are solely responsible for determining, reporting, and remitting all applicable tax obligations related to AMA token acquisition, holding, use, and disposal. Amadeus is an open-source project distributed on an "as is" basis. To the maximum extent permitted by applicable law, the Amadeus team, its contributors, developers, advisors, and affiliates expressly disclaim all liability for any direct, indirect, incidental, consequential, or punitive losses or damages arising from or related to this litepaper, the Amadeus network, or AMA tokens, including but not limited to loss of value, theft, technical failures, security breaches, or regulatory actions. Prospective participants must conduct independent due diligence and possess sufficient knowledge, experience, and sophistication regarding blockchain technology and cryptographic tokens to properly evaluate the risks and merits of participation. By accessing this litepaper or acquiring AMA tokens, you expressly acknowledge that you have read and understood these risks, have conducted your own independent investigation, accept full responsibility for your decisions, and irrevocably agree not to hold the Amadeus team or any related parties liable for any losses, damages, or adverse consequences whatsoever.

Last Updated: December 2025

Acknowledgement of Contributors to the Litepaper

This litepaper represents a collaborative effort from team members of the Amadeus team, advisors and community members. The following key members helped author this litepaper:

Ivan Svirid (Vans), Valentyn Faychuk (Valy), Leslie Kivit (Talentnodes), Laura Beniguel (LauraX), Dr. Kody Law, Dr. Andrea Corceone, Aaron Law, Tatiana Hernandez, SnowFlax, Benkei, Aurelius, Hiroshima, Kengan, Banto and TypeError.