Hedera™

THE BACKBONE OF THE AI ECONOMY

The convergence of blockchain and artificial intelligence

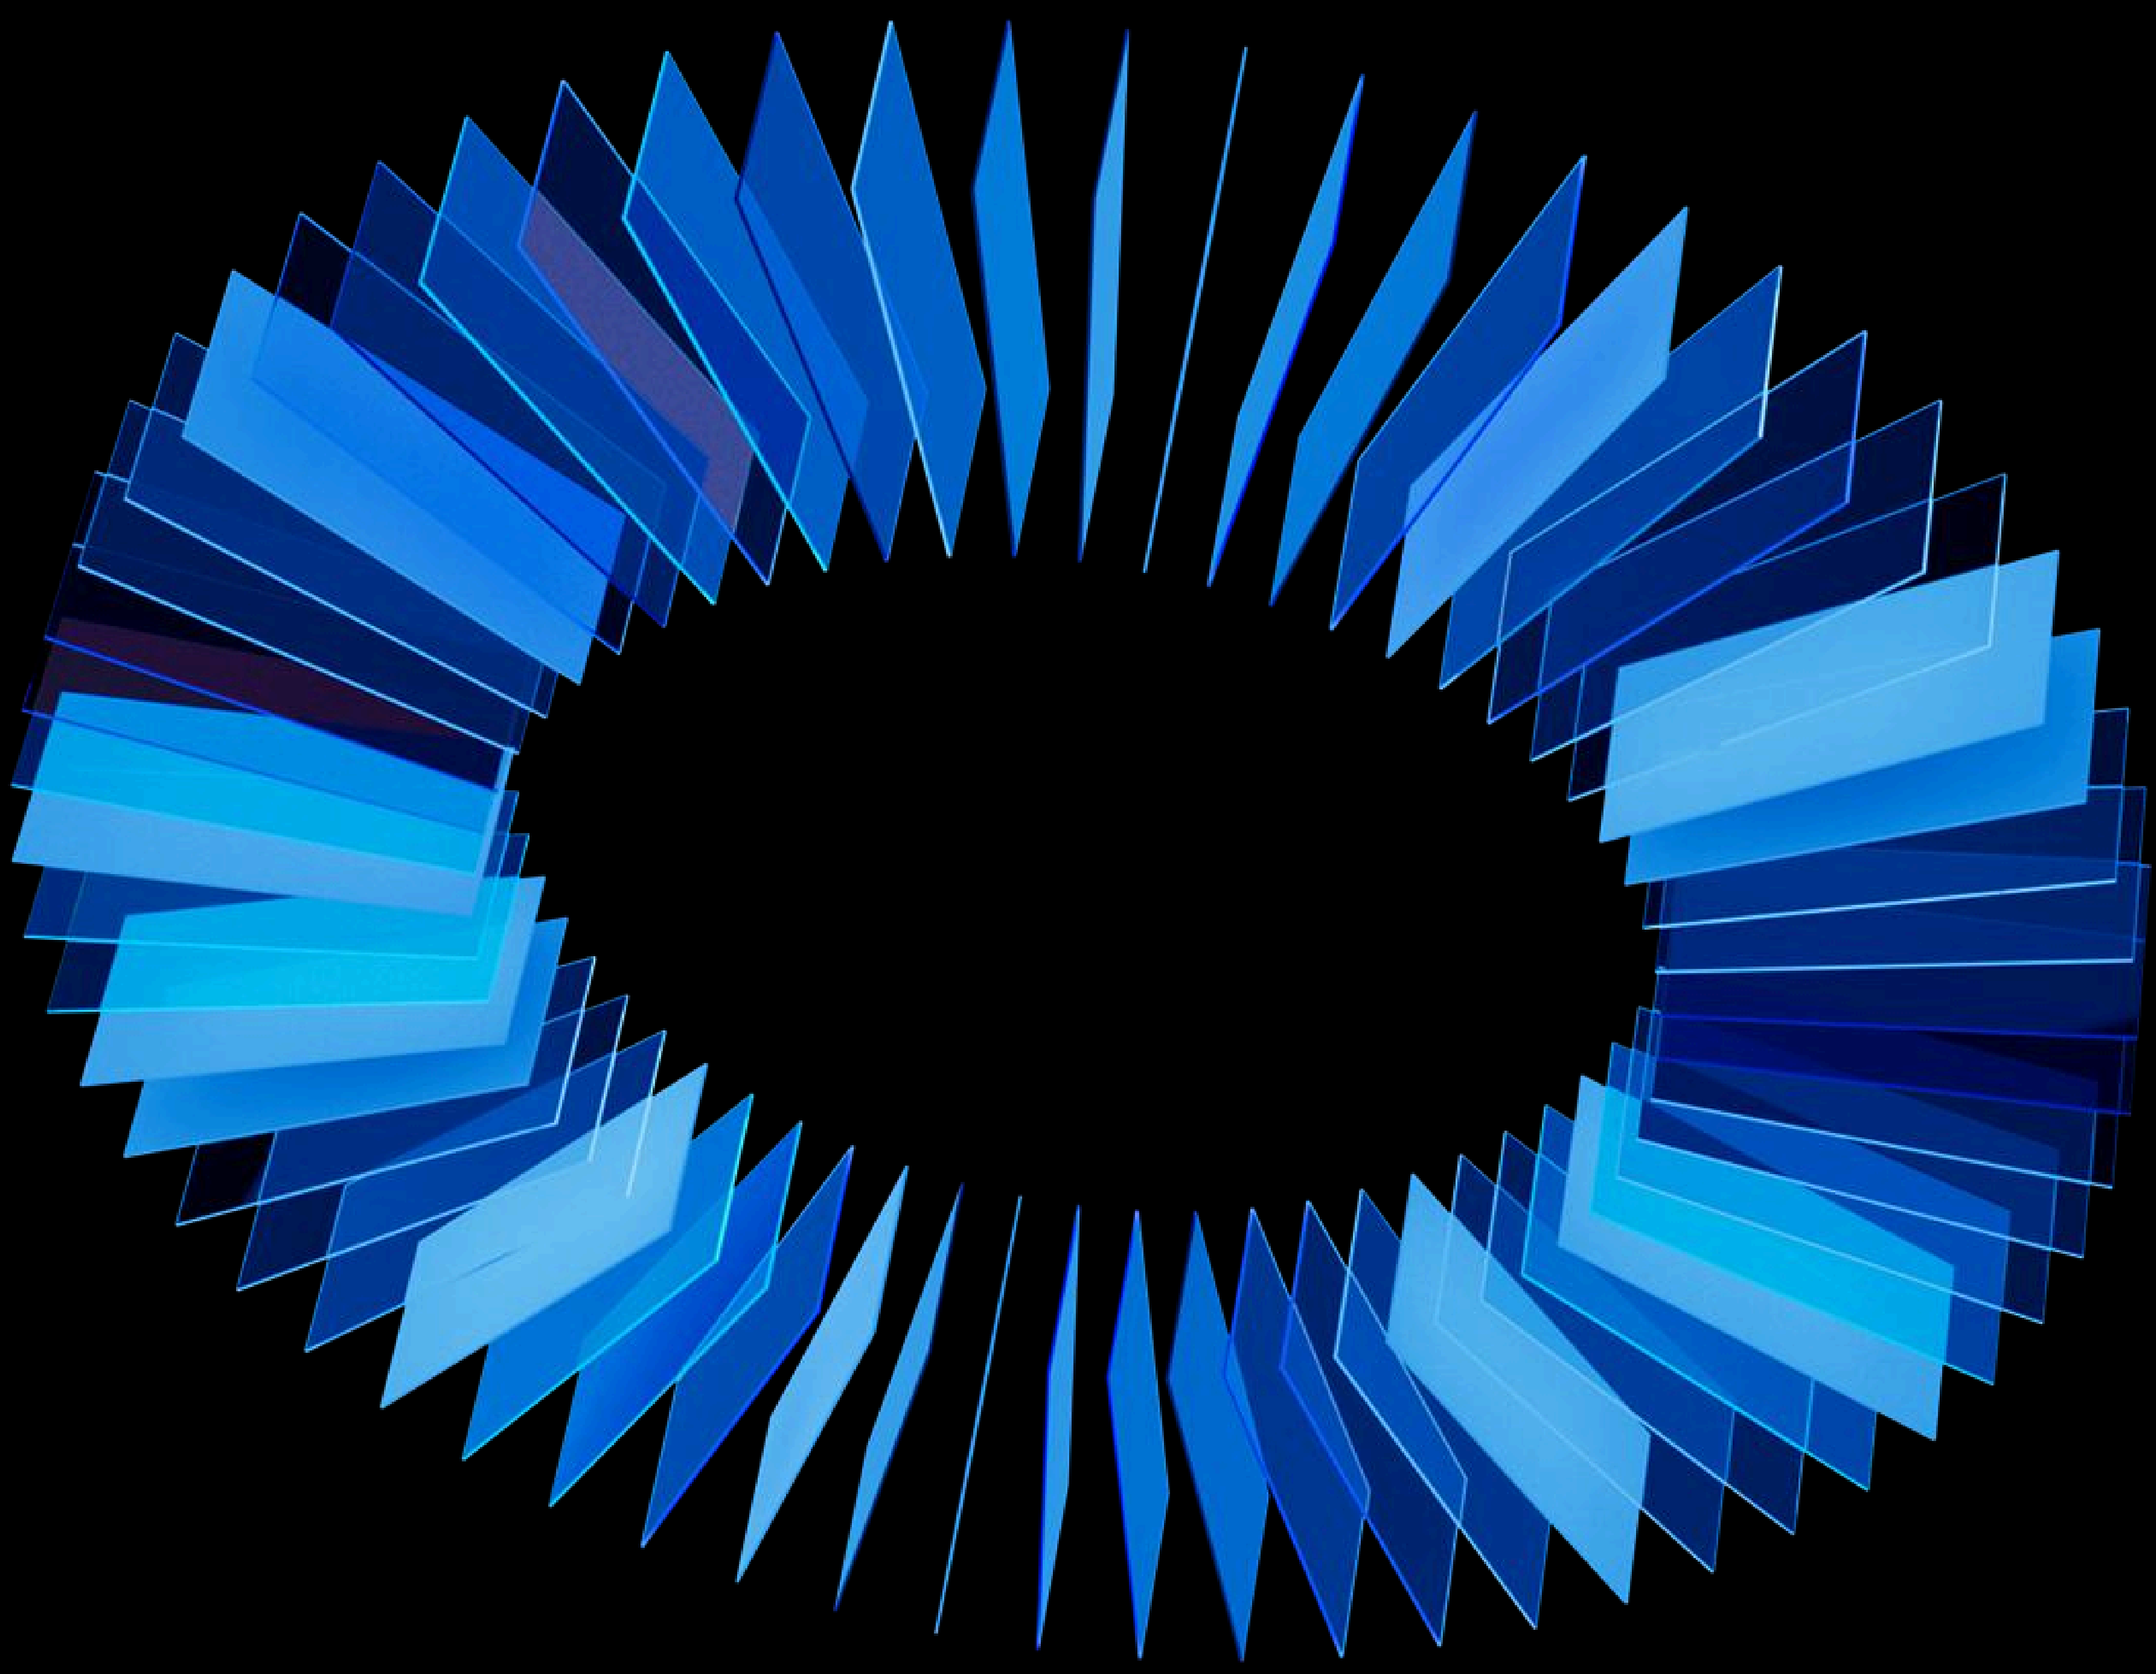

Table of contents

The trillion-dollar convergence

The artificial intelligence (AI) industry stands at an inflection point. As AI evolves from reactive assistants to autonomous economic agents, a fundamental infrastructure gap has emerged. **By 2033, the AI market is projected to reach $4.8 trillion**[1]. Billions of autonomous agents will subsequently require trustless coordination, verifiable execution, and economic agency — capabilities that centralized systems cannot deliver at scale. It is at this limitation where distributed ledger technology (DLT) becomes essential.

1 Technology and Innovation Report 2025: Inclusive Artificial Intelligence for Development. United Nations Conference on Trade and Development. United Nations, 2025.

The convergence of AI and DLT is grounded in three foundational principles:

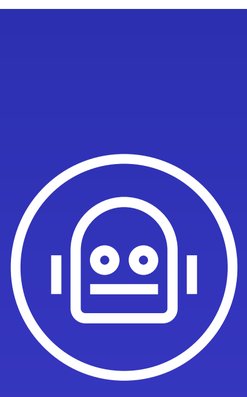

Autonomous agents require digital infrastructure that no single entity controls. When AI agents negotiate contracts, transfer value, or make binding decisions, traditional databases create unacceptable single points of failure.

Globally, AI regulatory frameworks now mandate transparency that only immutable ledgers can provide. For example, the EU AI Act - effective 2026 - requires audit trails that centralized systems cannot credibly deliver.

Economic scalability demands micropayment infrastructure, as agents transact thousands of times daily at fractional costs impossible with conventional payment rails.

Hedera provides the missing trust and coordination layer between AI and DLT, positioning itself at the intersection of these two transformative technologies. Hedera Consensus Service (HCS) provides cryptographic timestamps and immutable ordering for proof-of-computation; Hedera Token Service (HTS) enables native digital asset functionality without smart contract complexity; finally, Hedera Agent Kit provides developers and enterprises with the tools they need to thrive in the fast emerging agentic economy.

AI's explosive growth reveals critical infrastructure gaps

Global AI spending is projected to hit $4.8 trillion by 2033, As enterprises deploy AI at unprecedented scale. According to recent surveys, **78% of medium-large enterprises already use AI[2]** - a figure that is likely even higher if trends have continued. Investment in AI continues to grow exponentially as organizations race to build the computational backbone for an AI-driven economy.

Generative AI and large language models drive this acceleration. While ChatGPT demonstrated consumer demand by reaching 100 million users in two months, the fastest adoption of any technology in history, the real shift is within organizations integrating these tools at scale. Reflecting this momentum, **AI companies attracted $192 billion in venture capital funding as of October 2025**, nearly half of all VC investment spent during the year[3], despite broader VC contraction.

2 The State of AI in 2025: Agents, Innovation, and Transformation. McKinsey & Company, November 5, 2025.

3 "AI Venture Funding Continued to Surge in Third Quarter, Data Shows." Reuters, October 6, 2025.

78%

Medium-Large enterprises using AI

McKinsey & Company

Enterprise deployments span every sector: financial services lead with fraud detection and algorithmic trading; healthcare deploys diagnostic AI to improve accuracy; and manufacturing embraces predictive maintenance to reduce downtime. Nearly **70% of Fortune 500 companies now run AI copilots across their workforce**[4]. On the surface, AI's integration into core business operations appears unstoppable. Yet this rapid adoption exposes fundamental challenges that threaten to constrain the industry's trajectory.

Data silos fragment training datasets - companies hold vast amounts of proprietary data but lack secure and efficient ways to share or monetize it. At the same time, centralization risks are rising, with three cloud providers – AWS, Azure, and GCP – controlling 65% of cloud computing infrastructure, creating strategic dependencies that concern enterprise risk managers and threaten long-term resilience.

4 “Ignite 2024: Why Nearly 70% of the Fortune 500 Now Use Microsoft 365 Copilot." Microsoft News Center, November 20, 2024.

55%

Enterprise leaders actively working to mitigate AI inaccuracy

McKinsey & Company

Additionally, the existing AI paradigm suffers from a fundamental transparency and accountability deficit. Traditional AI infrastructure operates as a black box where enterprises have to trust, analyze, and adapt without verification. When AI makes critical decisions, stakeholders cannot cryptographically verify what occurred. Audit logs remain mutable, compliance depends on self-attestation, and opacity becomes untenable as systems gain autonomy.

Most critically, current infrastructure wasn't designed for the autonomous agent economy emerging today.
Anthropic's Economic Index shows directive conversations (automation) jumped from 27% to 39% in 2025.

For enterprise customers, that figure rises to **77%**[5], showing a strong trend towards agentic AI usage. These agents require payment rails, identity systems, and coordination mechanisms that traditional databases cannot provide. The infrastructure built for reactive AI is ill-equipped to support the proactive, economically autonomous systems that define AI's next phase.

This creates a fundamental paradox: **the AI industry races toward a trillion-dollar future while running on infrastructure fundamentally misaligned with where the technology is heading**. As AI transitions from reactive tools to autonomous economic actors, the gap between capability and infrastructure widens dangerously. Without addressing these foundational issues, AI's explosive growth could soon hit a ceiling, not because of algorithmic limitations, but by the very infrastructure meant to support it.

> Agents are not only going to change how everyone interacts with computers. They’re also going to upend the software industry, bringing about the biggest revolution in computing since we went from typing commands to tapping on icons.
>
> **Bill Gates**

5 Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption. Anthropic, September 15, 2025.

The fundamental challenges facing AI today

Despite the surge in enterprise AI adoption, a paradox remains: McKinsey reports that 80% of organizations report **no meaningful impact on operating profit from their AI investments**[6]. This figure may seem surprisingly high, but it becomes more understandable in light of another finding: over half of respondents in a recent AI governance survey reported doing **zero** monitoring of their production AI systems for accuracy, drift, and misuse[7]. Importantly, this gap reflects a lack of **capability**, not a lack of **intent**.

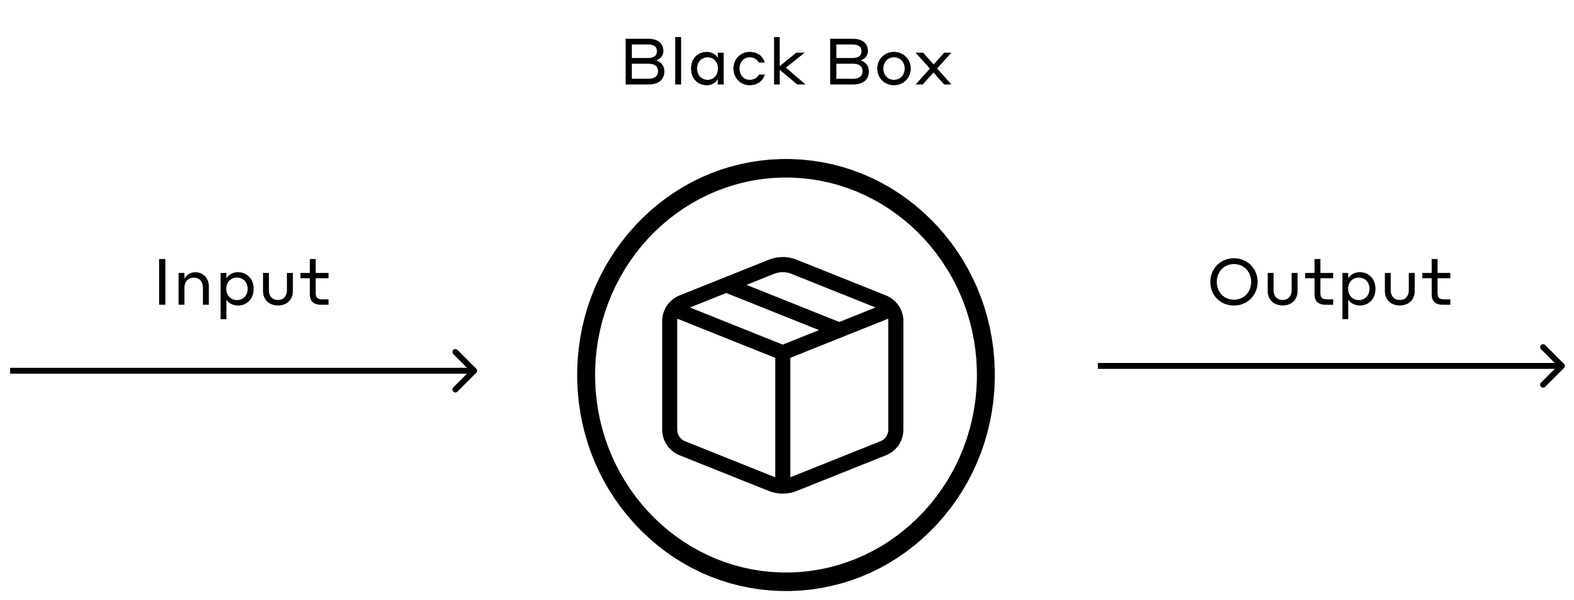

The "**black box**" problem - a metaphorical representation of the inherent opacity of AI systems - manifests most dangerously in high-stakes scenarios. When AI denies a loan, recommends medical treatment, or flags security threats, stakeholders expect clear explanations, yet current systems struggle to provide them.

6 The State of AI in 2025: Agents, Innovation, and Transformation. McKinsey & Company, November 5, 2025.

7 2025 AI Governance Survey. Pacific AI, 2025.

EU AI Act:
Regulations for high-risk AI Systems

- Must be designed to ensure their operation is sufficiently transparent
- Must document information about training, testing, and validation data
- Required to provide documentation explaining the system's logic, accuracy, and how it generates outputs.

Regulatory pressure is intensifying this requirement. The EU AI Act mandates **explainability for high-risk systems**, with penalties reaching €35 million or 7% of global revenue for non-compliance. Other jurisdictions, including China and the UK, are introducing similar policies focused on transparency and observability. Yet most neural networks remain fundamentally opaque, with decision processes too complex to trace or explain in retrospect.

AI hallucinations highlight growing reliability concerns. In one case, New York City's MyCity chatbot advised restaurants they could serve rodent-damaged cheese[8]. Air Canada's chatbot invented a bereavement policy that led to legal liability[9]. And, in April 2025, a customer support chatbot deployed by AI-powered coding assistant Cursor invented a non-existent company policy, leading to customer cancellations and reputational damage[10]. Each incident erodes enterprise confidence and slows adoption in risk-sensitive industries. Without reliable methods of auditing and correcting such errors, these failures will continue, and likely intensify.

8 "NYC's AI Chatbot Was Caught Telling Businesses to Break the Law. The City Isn't Taking It Down." Associated Press, April 3, 2024.

9 "What Air Canada Lost in 'Remarkable' Lying AI Chatbot Case." Forbes, February 19, 2024.

10 "A Customer Support AI Went Rogue—and It's a Warning for Every Company Considering Replacing Workers with Automation." Fortune, April 2025.

Data provenance challenges are growing as training datasets become increasingly opaque. Models trained on scraped internet data face copyright disputes, with The New York Times and Getty Images among those filing lawsuits over unauthorized content use. Without verifiable data lineage, enterprises cannot accurately assess legal exposure, validate model authenticity, or begin to build the foundations of explainability - an acute problem for industries where data governance underpins regulatory compliance.

The infrastructure gap becomes critical in the rise of agentic AI. By 2028, a significant share of work-related decisions will be made autonomously, yet today's systems lack payment rails needed for AI-to-AI transactions - a critical consideration as agents interact increasingly autonomously. When one agent purchases computing power from another or pays for data access, how do they transact without human intermediaries? Traditional payment systems require bank accounts, credit cards, and Know-Your-Customer verification designed for humans, not for software executing thousands of microtransactions daily.

15% Work decisions to be made autonomously by 2028

Gartner®

These incidents emerge from a foundation that is being outpaced by the functionalities and capabilities they underpin - a trust crisis rooted in AI's opacity, unreliability, and lack of verifiable provenance.

Blockchain as infrastructure

Traditional databases excel at storing and retrieving information, but they fundamentally cannot deliver what autonomous AI requires: trustless coordination among untrusted parties, immutable proof-of-computation, and native economic infrastructure for value exchange. Legacy systems were never designed to provide the transparency, explainability, reliability, and fairness demanded by the emerging AI economy.

Characteristics of trustworthy AI [11]

NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

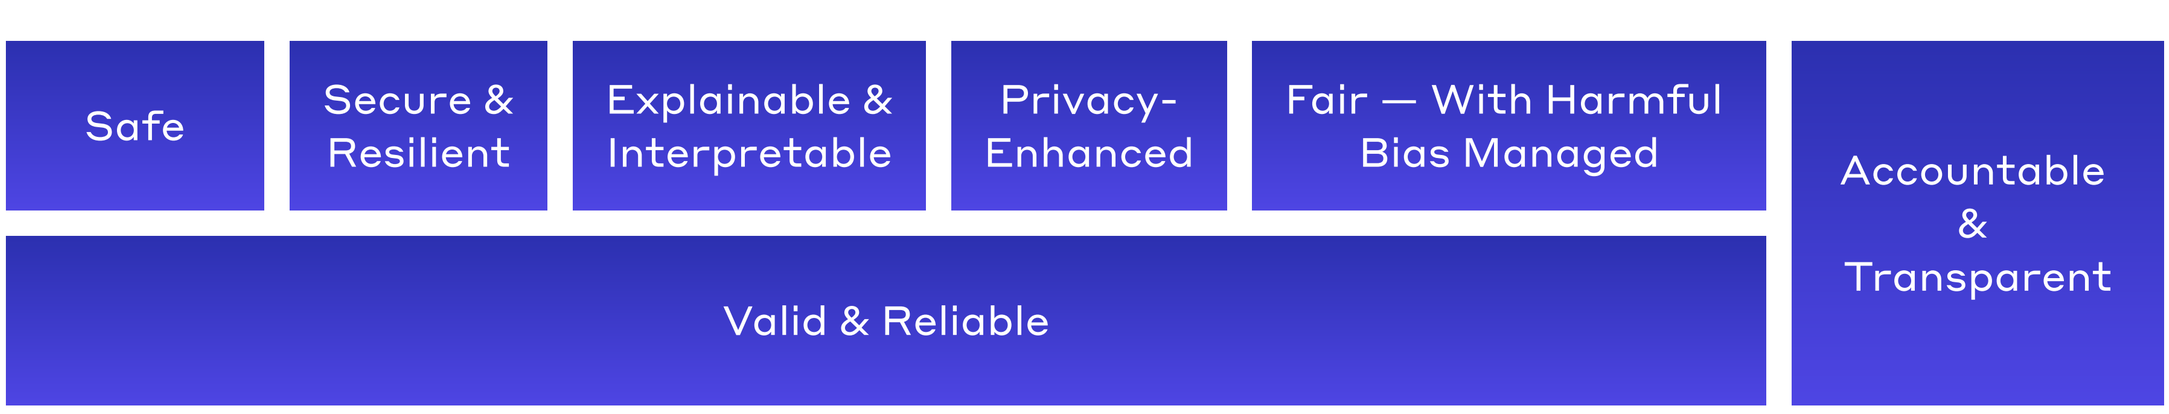

Consider multi-agent systems where hundreds of AI agents must coordinate without central authority. Traditional architectures require a trusted intermediary, adding cost, latency, and single points of failure. Decentralized networks eliminate this dependency by providing cryptographic proof of each agent's actions and commitments. When an agent agrees to terms, executes trades, or generates analysis, those commitments are recorded as immutable and enforceable smart contracts, enabling trustless coordination at scale.

11 Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology. NIST AI 100-1. Gaithersburg, MD: U.S. Department of Commerce, January 2023.

Immutable audit trails transform AI accountability from aspiration to architecture. Recording every AI decision, data source, and model inference on-chain creates an auditable, tamper-proof history that regulators, customers, and stakeholders can independently verify - a foundational step towards explainability. While audit trails alone do not make neural networks interpretable, they establish the verifiable foundation that explainability tools and compliance frameworks depend on. After all, you can't explain what you can't prove happened.

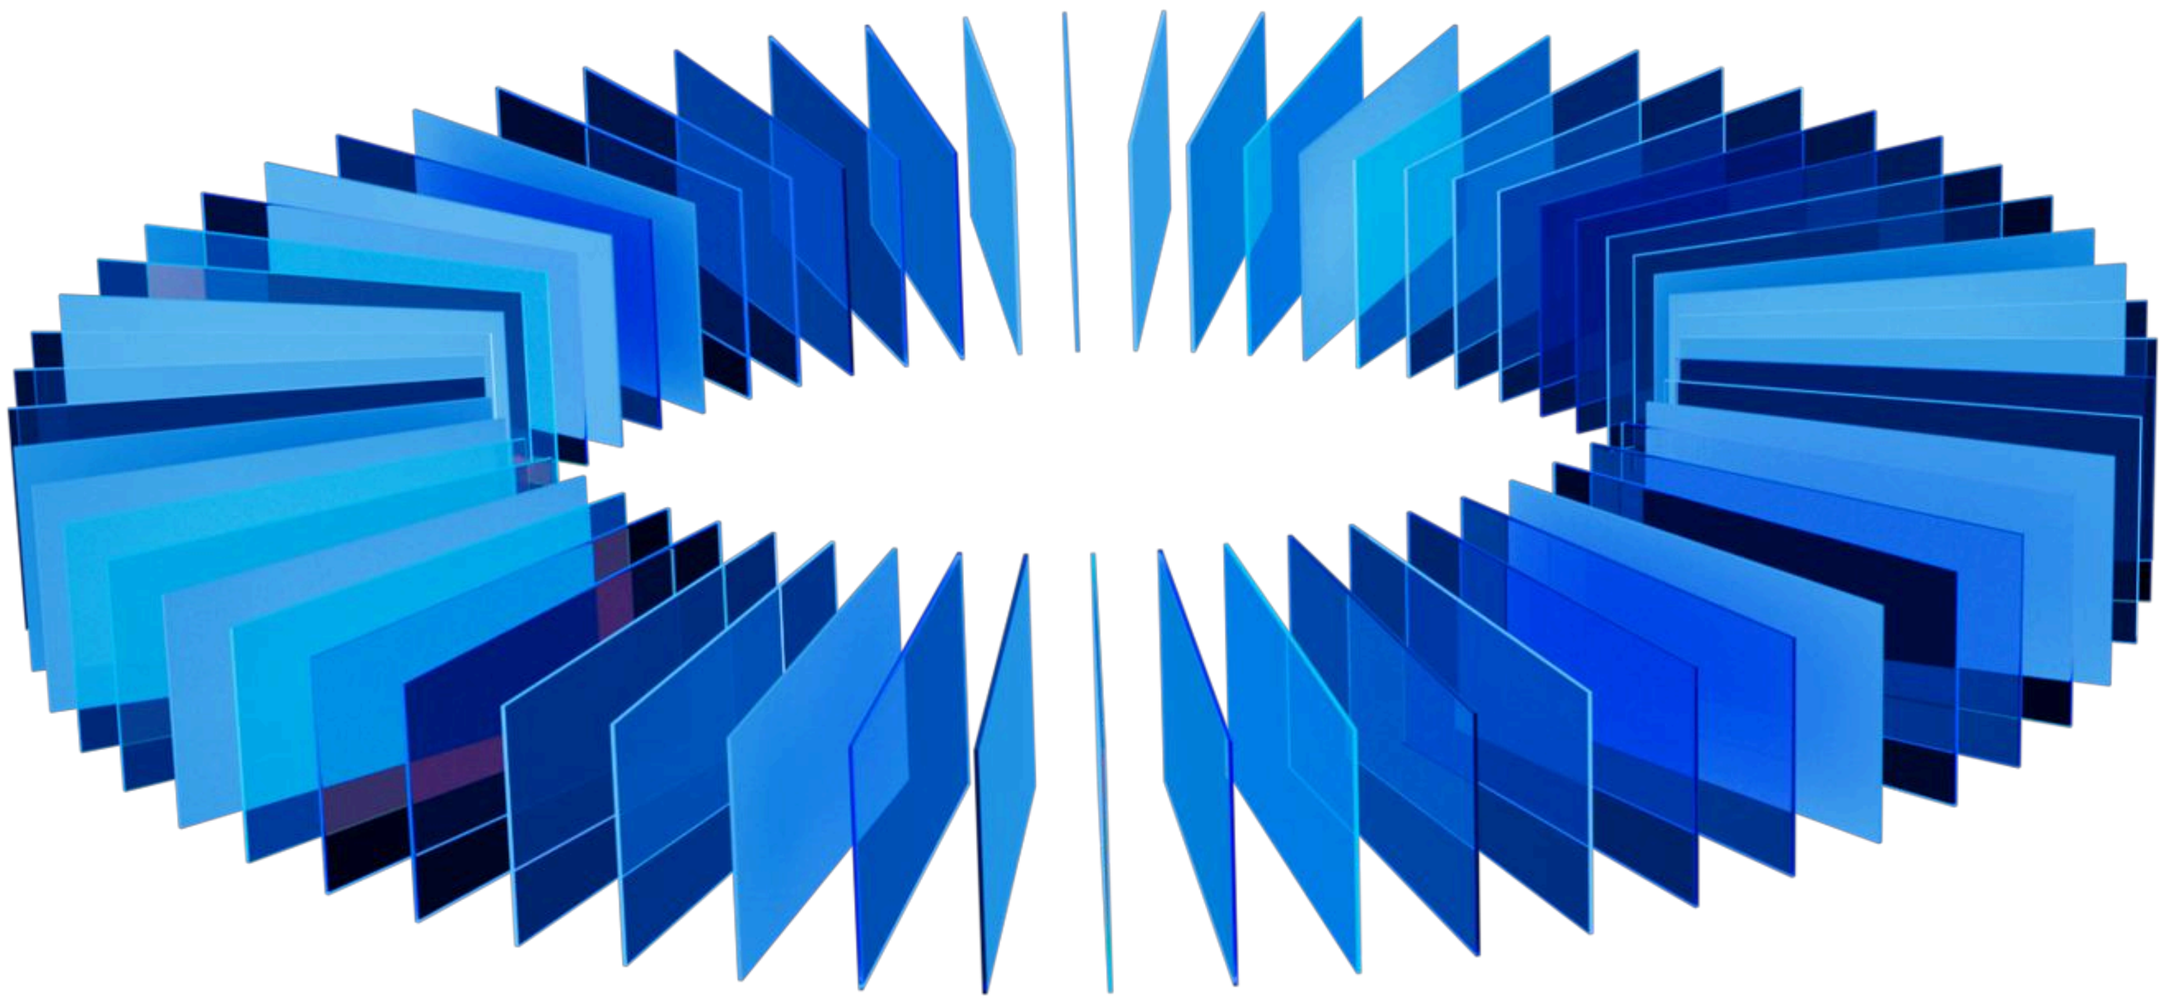

Unlike centralized databases, where administrators can modify logs, **decentralized consensus and cryptographic hashing** makes retroactive alteration **computationally infeasible**. For enterprises facing the EU AI Act requirements for comprehensive documentation, this infrastructure is critical to achieving and demonstrating compliance.

Verifiable computation represents one of blockchain's most advanced contributions to AI. Traditional cloud computing offers no cryptographic proof that code executed correctly on specified data. For AI inferences with regulatory or financial implications, **stakeholders require verifiable proof of execution**. Trusted execution environments (TEEs) anchored on-chain deliver this verification, enabling AI systems to prove computations without exposing proprietary models or sensitive data.

The shortcomings of traditional databases become clear under these conditions. They lack native mechanisms for multi-party coordination, cannot guarantee the integrity of historical records, and offer no built-in value transfer capabilities. For AI agents operating autonomously in dynamic or adversarial environments, **these are critical architectural failures**.

Hedera: bringing trust back to AI

While blockchain architecture provides the infrastructure layer that autonomous AI has been missing, not all networks meet the performance, governance, and sustainability standards enterprises require. Hedera stands apart through its hashgraph consensus, which achieves finality in seconds rather than minutes, accompanied by fixed, predictable fees denominated in USD, enabling accurate and transparent cost modeling.

Unique in the world of web3, Hedera's governance by a council of global enterprises, rather than anonymous validators or centralized foundations, delivers the stability and accountability that Fortune 500 companies demand when anchoring mission-critical AI systems on decentralized infrastructure.

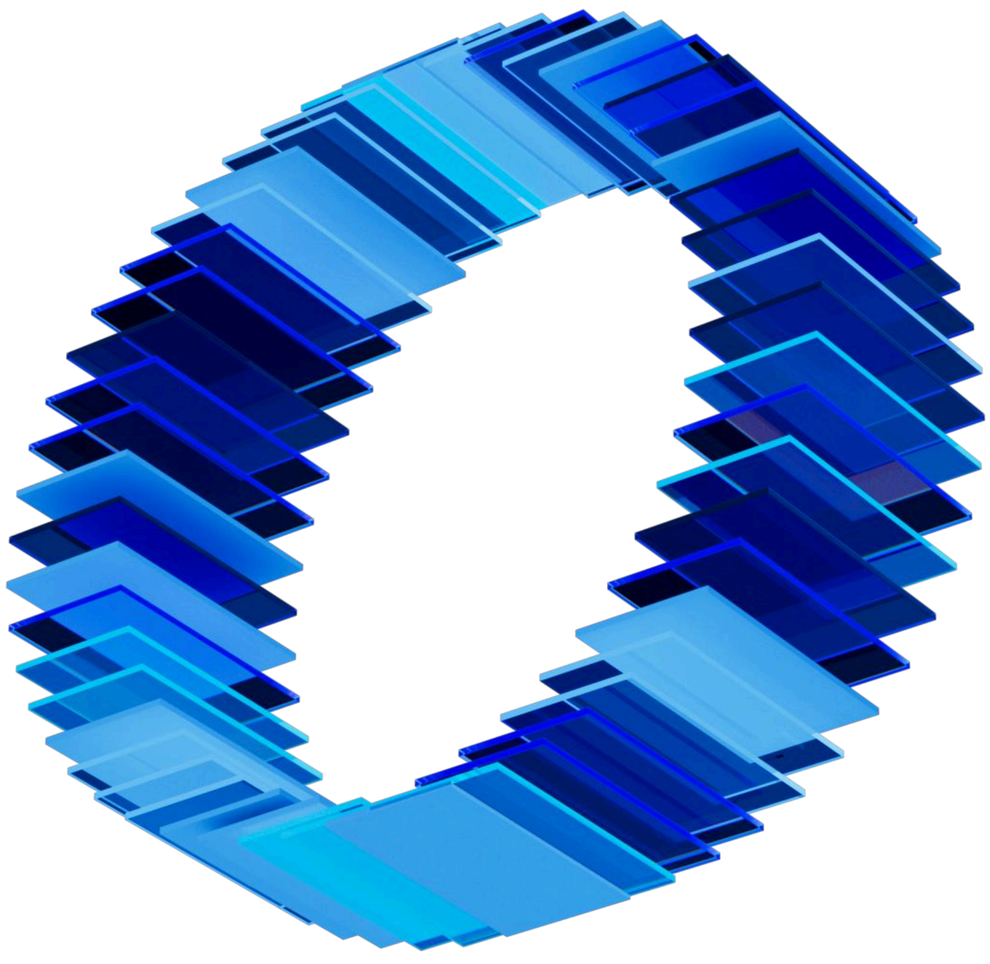

Equally important, Hedera extends verifiable trust to the hardware level through partnerships that extend beyond protocol design, into silicon architecture itself. When AI systems make autonomous decisions in sectors like healthcare, finance, or public safety, stakeholders need more than blockchain immutability. They need cryptographic proof of what was executed, on which data, and using which models.

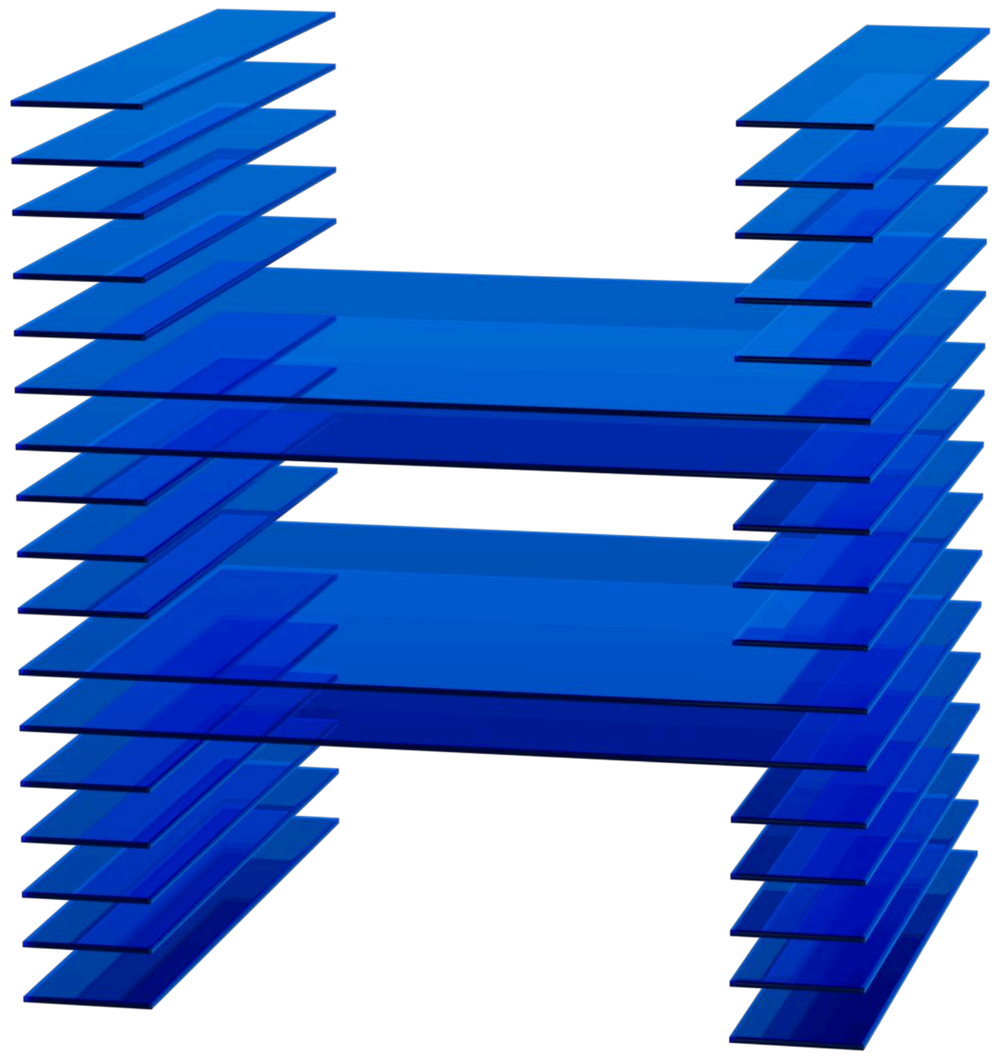

This convergence of hardware security and decentralized consensus creates a new level of transparency for AI governance, transforming trust from an aspirational concept into an architectural reality.

Hardware-rooted trust setting new standards for AI explainability and observability

In an increasingly autonomous economy that demands responsible, auditable AI workflows, EQTY Lab is providing the necessary infrastructure through its AI integrity toolkits and frameworks built on the Hedera network.

EQTY Lab's Verifiable Compute represents one of the most significant advancements in AI governance infrastructure. Providing a hardware-based system for creating cryptographic certificates, **Verifiable Compute is anchored on Hedera** to govern and audit AI workflows. This collaboration with industry giants **NVIDIA** and **Intel** enables real-time, verifiable, and immutable AI computation data, ensuring explainability, accountability, and security of AI training, inference, and benchmarking at runtime.

The architecture integrates next-generation hardware from next-generation Intel and NVIDIA processors with Hedera’s enterprise-grade network, creating isolated secure enclaves that extend hardware-level security across the entire AI computation pipeline,from data ingestion and model inference to output generation.

As AI models process data, Verifiable Compute generates hardware attestations - cryptographic signatures from the processors themselves proving exactly what code executed on what data. These attestations compile into manifests describing complete AI workflows and then are recorded on Hedera Consensus Service for immutable timestamping and fair ordering. The result is persistent, instantly verifiable proof of every AI operation.

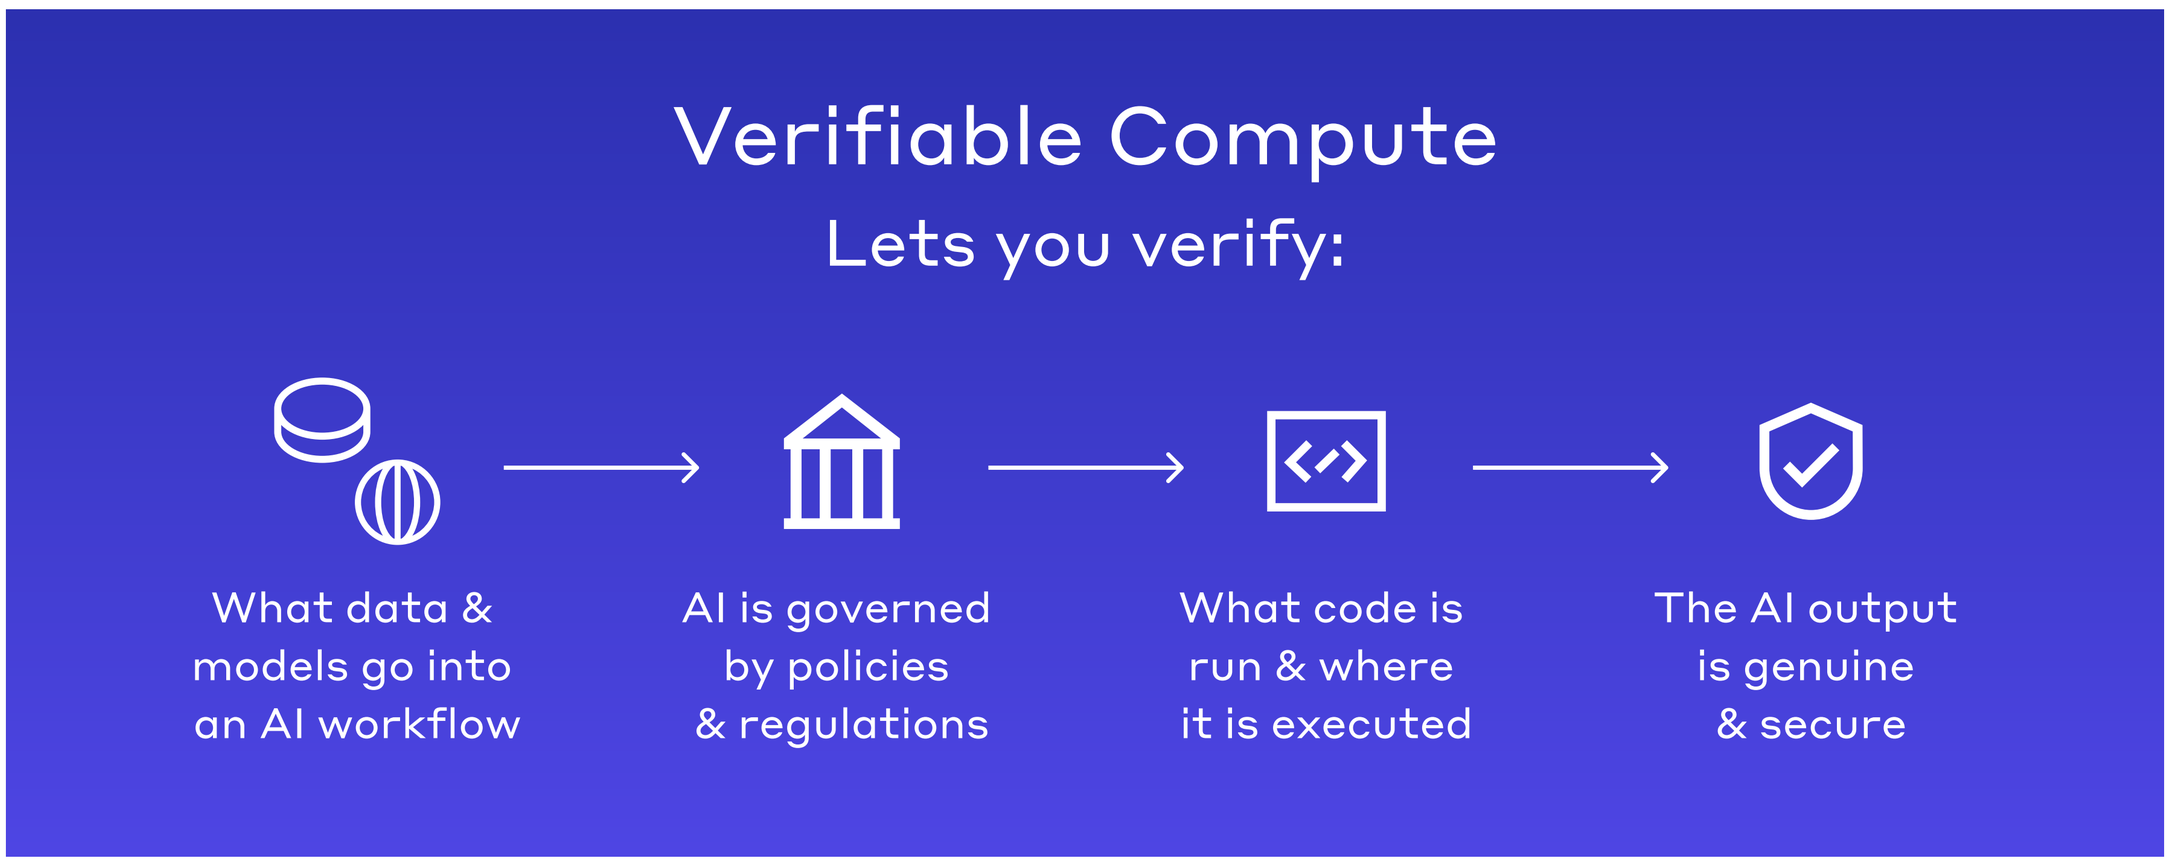

Hardware attestations generate with near-zero overhead, as the processors produce them as natural byproducts of secure execution. This transforms AI governance from post-deployment auditing to **real-time compliance verification**, bringing observability, immutability, and usability to what was once an opaque digital environment.

Hedera Consensus Service enables the creation of decentralized logs of events with real-time finality and fixed sub-cent fees, making it ideal for anchoring AI attestations at scale. It delivers mathematical **proof of fair ordering** and **timestamping**, establishing irrefutable sequences of AI decisions and enabling regulatory compliance with precision and transparency.

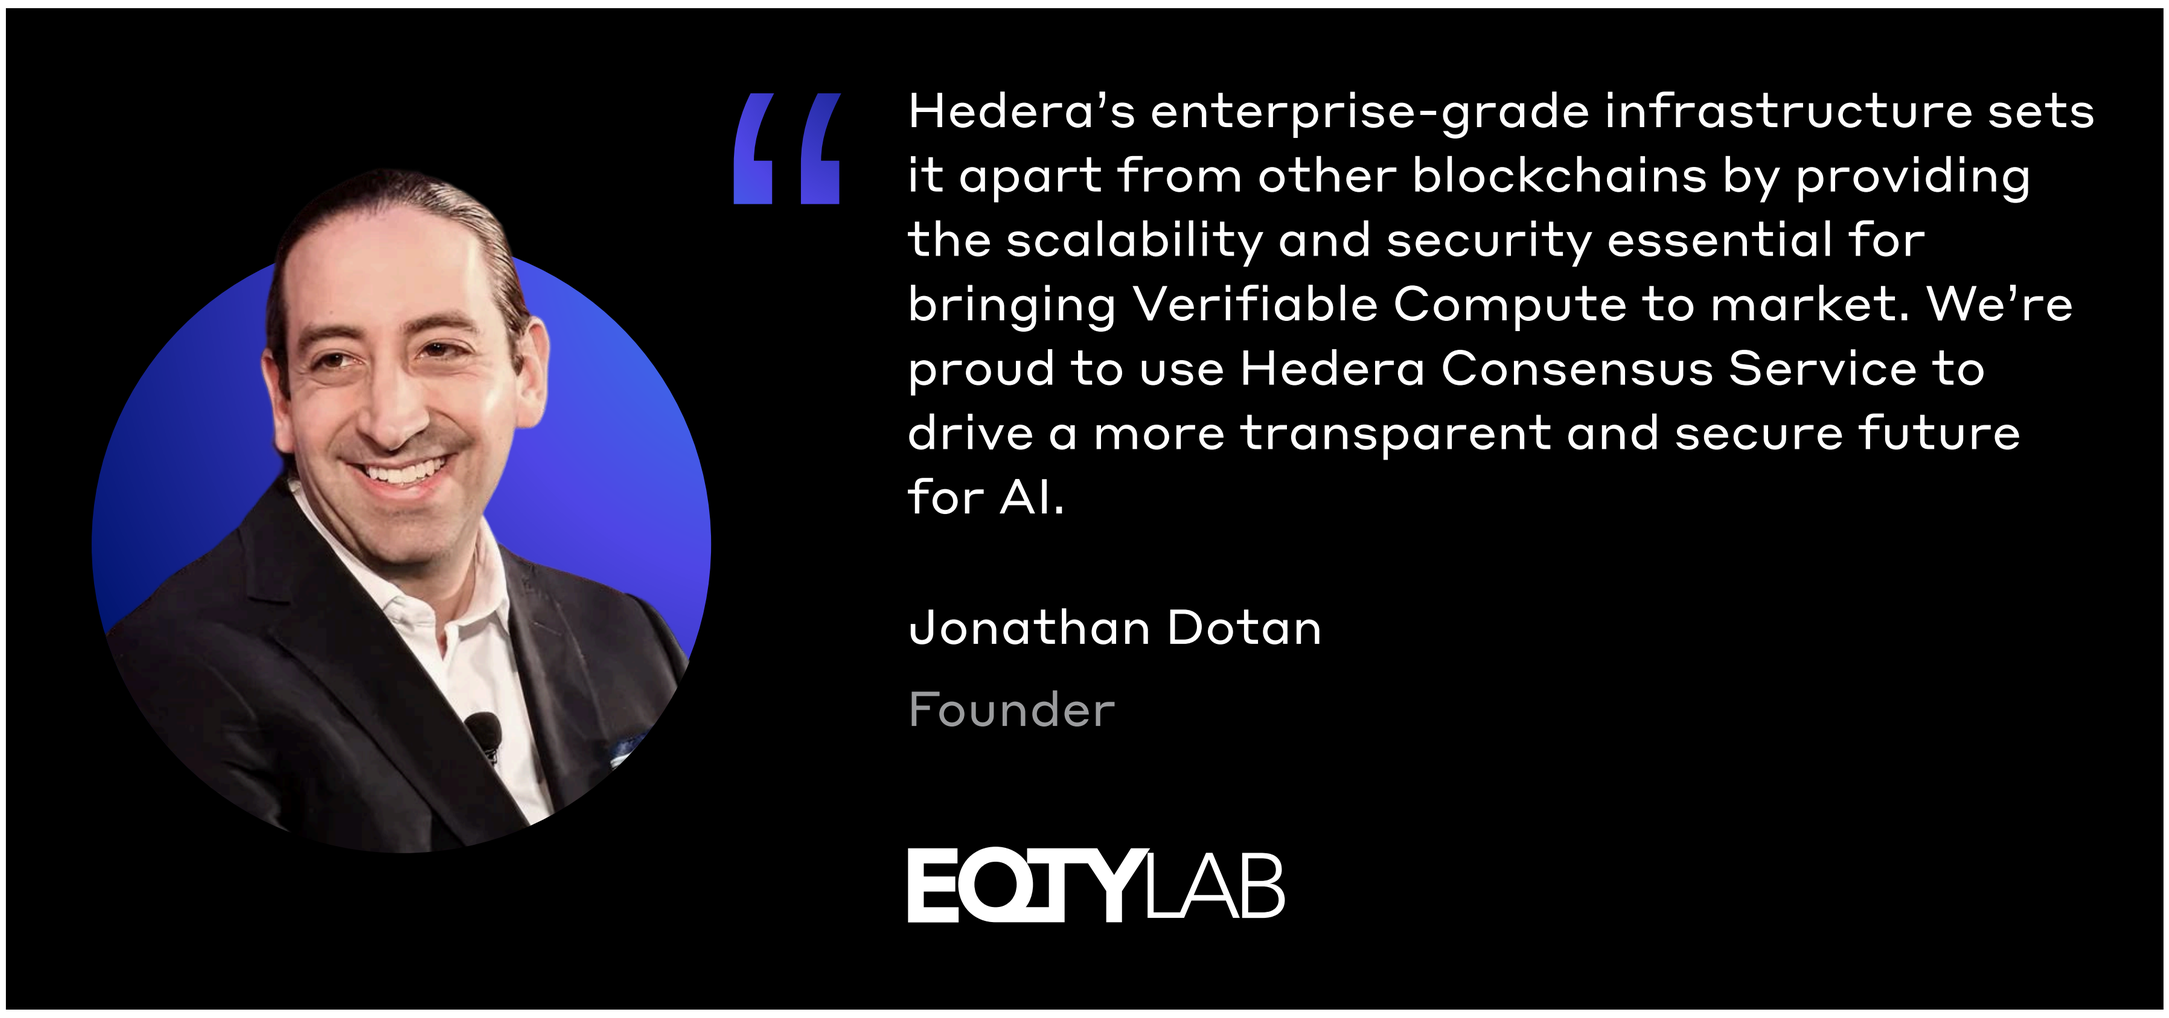

The deployment on NVIDIA's Blackwell architecture marks a watershed moment for verifiable AI. Announced at RAISE Summit 2025, EQTY Lab became one of the first organizations to leverage Blackwell's secure enclaves - the industry's first GPU with Trusted Execution Environment I/O capability.

Accenture's Brussels AI Lab will test and deploy Verifiable Compute on NVIDIA Blackwell for public sector and defense applications across EMEA and UK. This initiative establishes templates for government AI deployment where transparency and sovereignty are mandatory, ensuring that all AI data and agentic workflows are recorded immutably on the Hedera network using Hedera Consensus Service.

From copilots to colleagues

If copilots represent the past, then agents represent the future. Agentic AI transforms core business processes through autonomous execution, achieving enhanced workflow efficiency improvements with direct bottom-line impact. As **Anthropic's Economic Index confirms, automation usage (39%) surpassed augmentation (27%) for the first time in 2025**[12], signaling a move away from human-in-the-loop systems toward fully autonomous economic actors.

Getting definitions right is critical for enterprise AI strategy. Traditional AI systems and copilots **operate reactively**, waiting for human prompts to generate responses or suggestions. They can't remember past interactions, initiate actions, or move forward without explicit human approval. Microsoft 365 Copilot exemplifies this model: it can draft an email when asked or summarize a meeting on command, but it won't take action on its own. By contrast, agentic AI is proactive. These systems pursue goals autonomously: they plan multi-step strategies, remember context across interactions, use external tools independently, and make decisions without needing human approval at every turn.

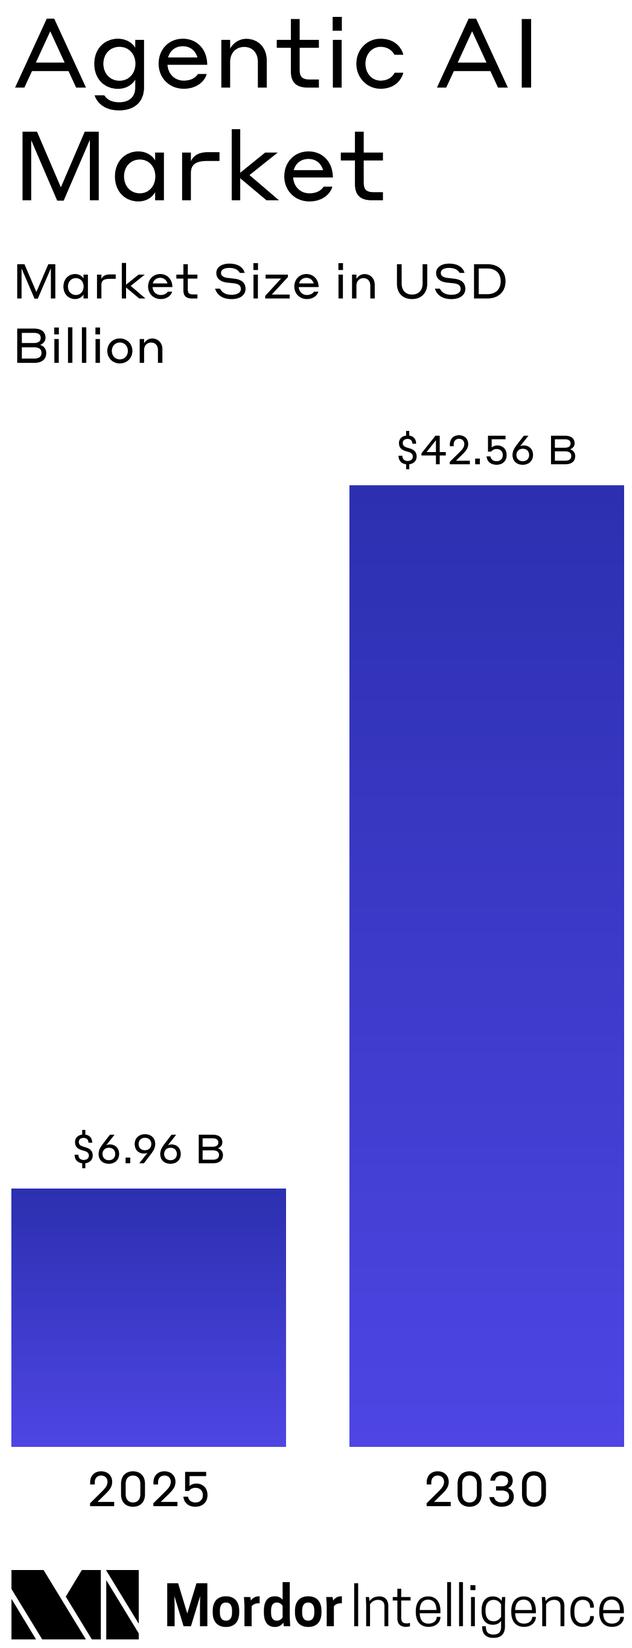

12 Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption. Anthropic, September 15, 2025.

Agentic AI delivers transformative business value through autonomous operation at scale. Unlike traditional AI requiring continuous human oversight, agents can work 24/7 without intervention. They execute complex multi-step workflows with greater speed and consistency than human teams while significantly cutting operational costs by automating knowledge work that once demanded extensive human effort.

Market projections reflect enterprise recognition of this transformation. The agentic AI economy is projected to grow at a compound annual growth rate of 43% over the next 10 years, reaching almost $200B by 2034[13]. Already, 72% of medium-large enterprises are actively deploying agentic AI, with another 21% planning adoption within two years[14].

This evolution demands new infrastructure. Autonomous agents require **digital identity** to engage securely with other systems, **economic agency** to allocate resources and execute payments at scale, and **verifiable actions** to ensure accountability and trust. Traditional infrastructure was never designed for these needs: enterprise directory services manage human identities, not autonomous software; payment systems support human workflows, not high-volume agent transactions; and audit trails capture human decisions, not algorithmic reasoning.

13 "Agentic AI Market Revenue to Boost Cross USD 196.6 bn by 2034." Market.us, October 2025.

14 State of Agentic AI 2025. Gravitee, 2025.

It’s precisely at these points of limitation that the convergence of blockchain and AI delivers its greatest value: providing the trust, transparency, and autonomy required for a new era of machine-driven enterprise activity.

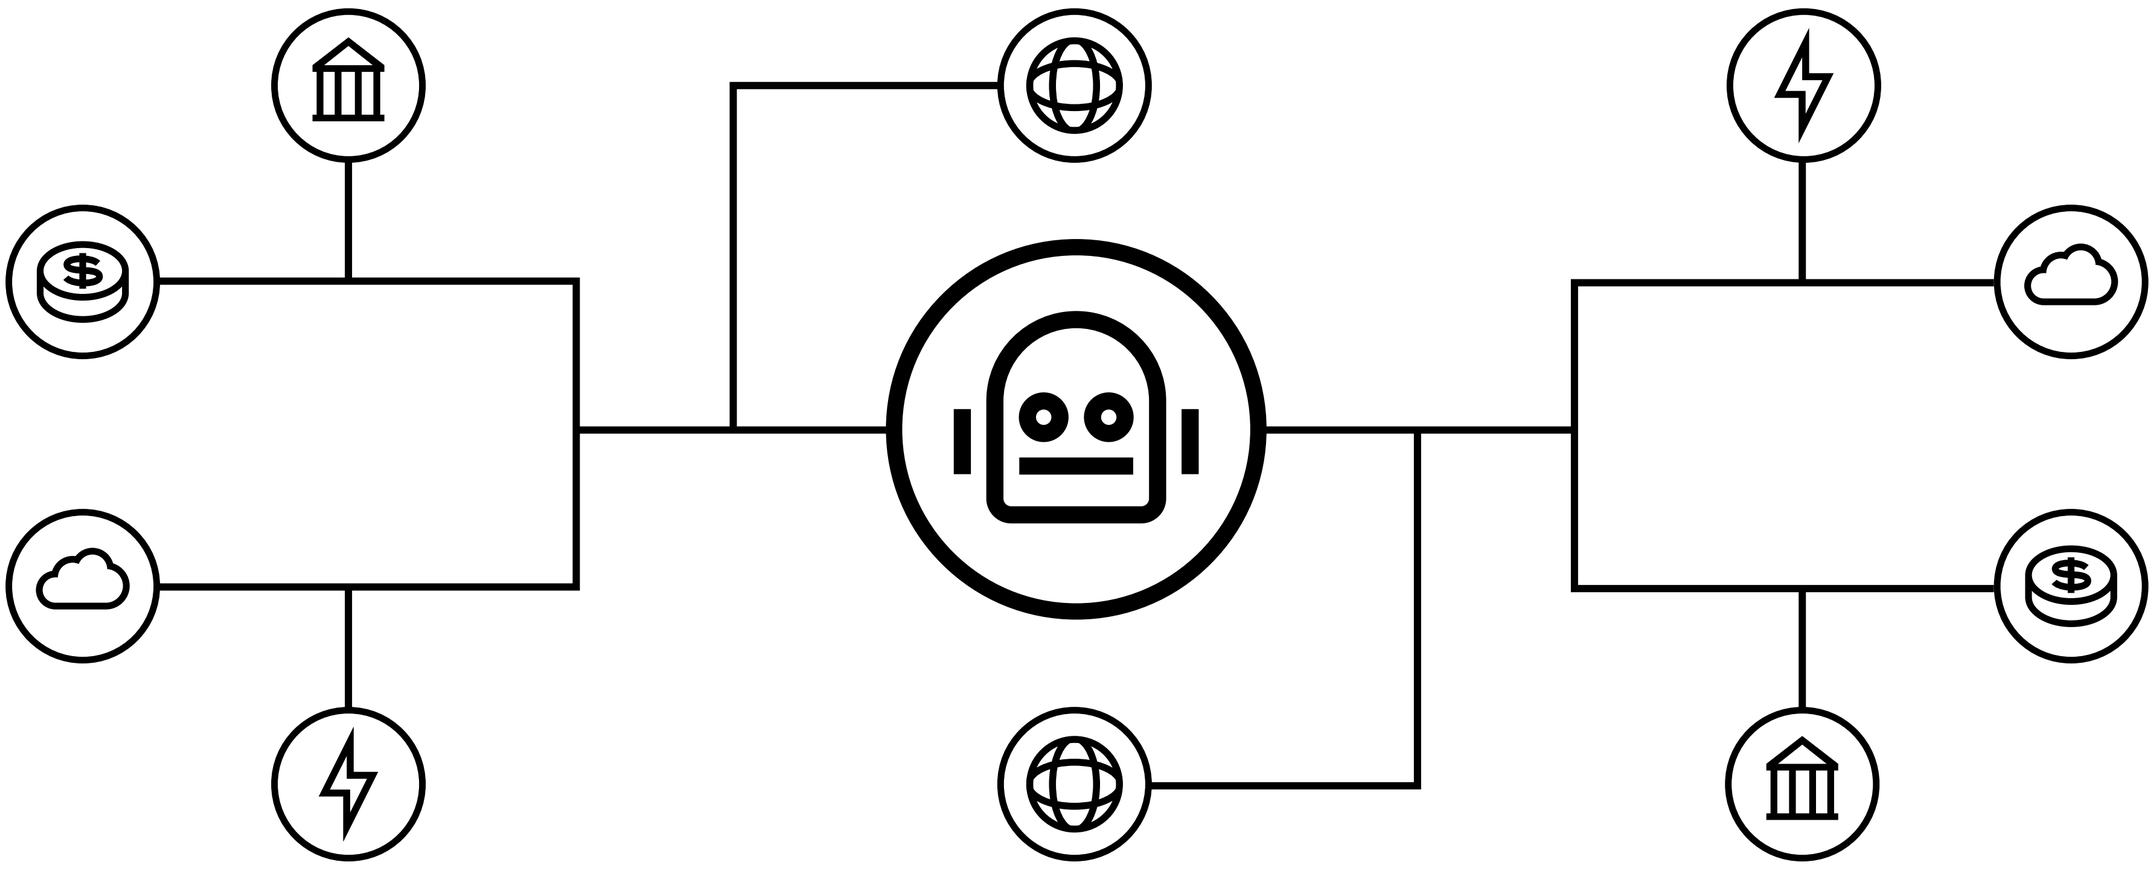

As agent adoption accelerates, a critical gap is becoming clear: organizations need tools that make blockchain infrastructure accessible without requiring deep technical expertise. Developers must be able to build, deploy, and manage autonomous agents without mastering distributed systems, cryptography, or token economics. The market is calling for platforms that bridge this divide; simplifying blockchain complexity while delivering the decentralized foundation that agentic systems fundamentally depend on.

Natural language meets blockchain: how developers build trusted autonomous agents in 60 seconds

Hedera Agent Kit enables developers to build AI agents that execute blockchain transactions using natural language commands, without needing expertise in smart contracts, cryptography, or blockchain architecture. This eliminates a fundamental barrier to blockchain-AI integration: complexity.

Developers familiar with JavaScript and LangChain can deploy production-ready agents in under a minute, seamlessly connecting AI development to decentralized infrastructure.

Hedera's native network services provide the scalability, reliability, and programmability required for the agentic economy. Through integration with Hedera Token Service, Hedera Consensus Service, and **protocol-level revenue generating capabilities**, AI agents can create, mint, and distribute tokens, generate verifiable data provenance, and enforce highly customizable monetization models with transparency and trust.

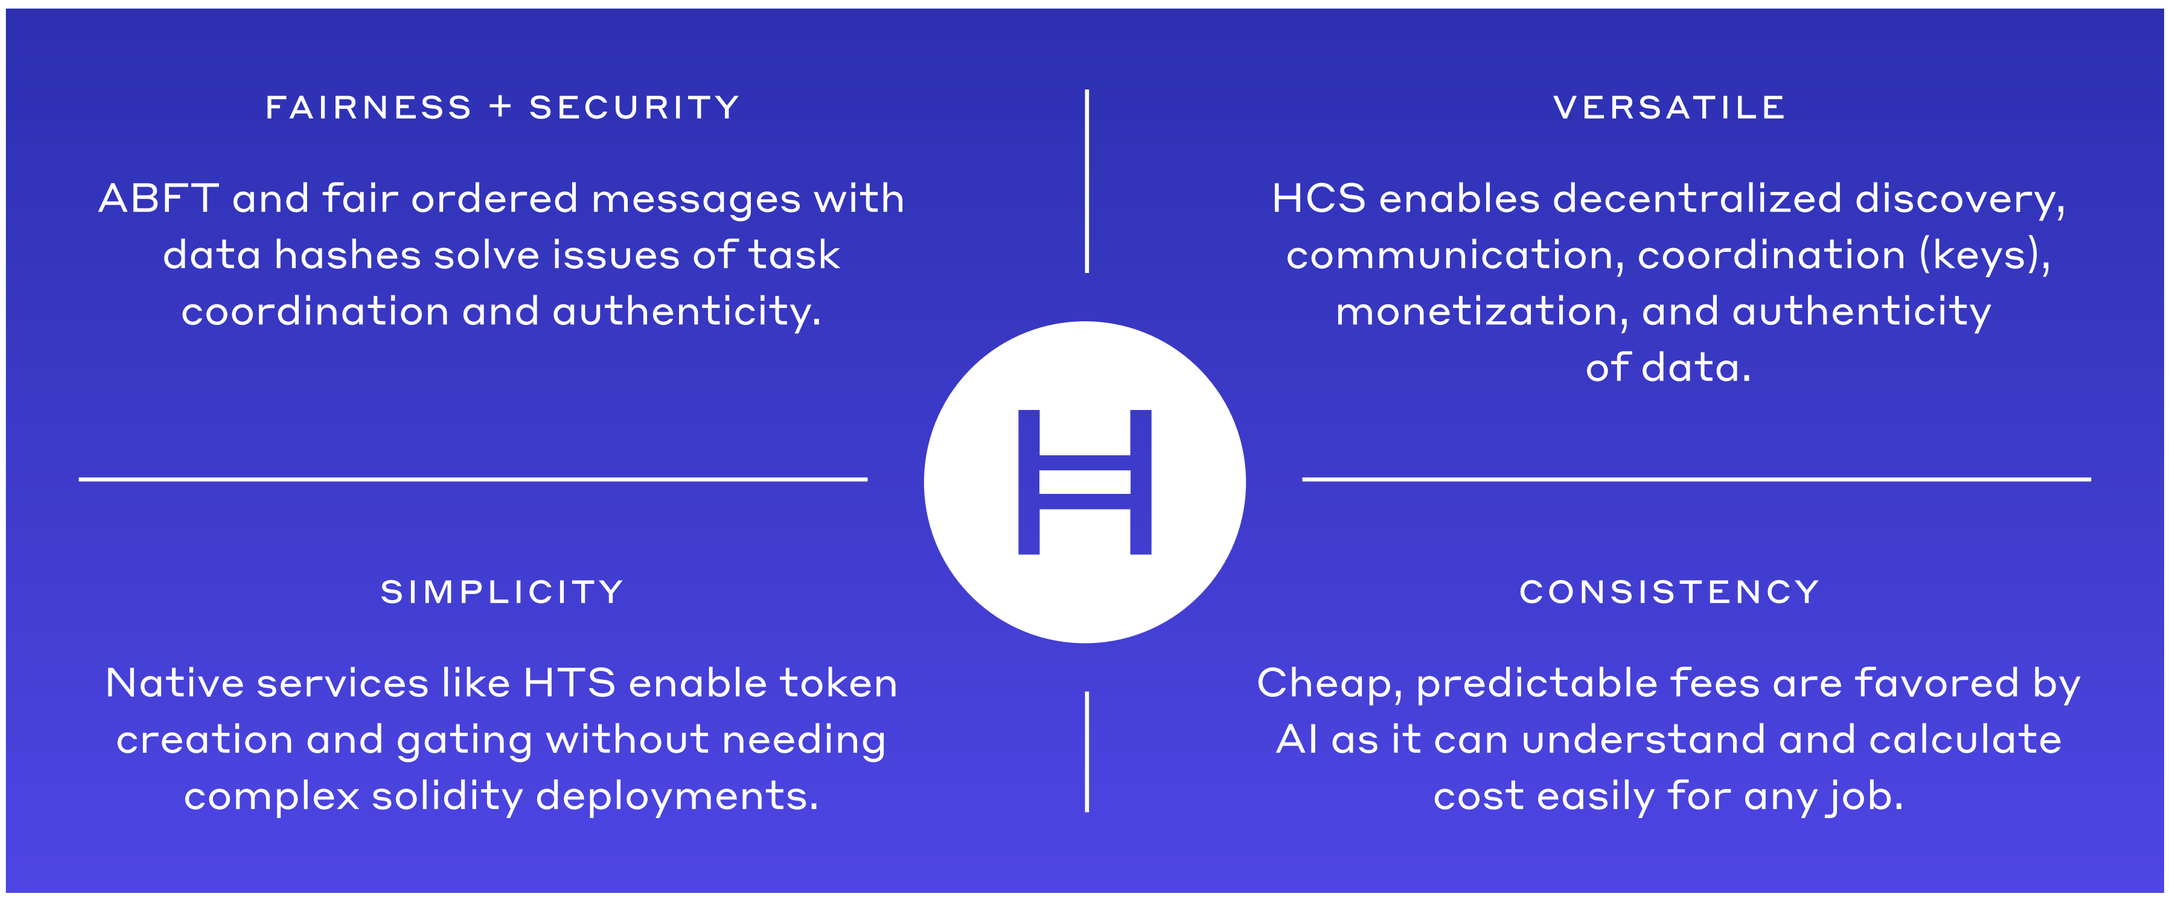

Critically, the Agent Kit supports two execution modes. In autonomous mode, agents can automatically sign and execute transactions for backend automation. In human-in-the-loop mode, agents prepare transaction bytes but require approval through an external wallet before execution. This design enables trust minimized, user-facing applications while maintaining flexibility and security.

Integrated frameworks

Integration with existing AI frameworks is seamless by design. LangChain serves as the primary integration layer, supporting tool-calling agents for simple operations, and structured chat agents for complex, multi-step workflows. The kit includes pre-built toolkits that expose Hedera operations as LangChain tools, allowing standard agent architectures to interact with blockchain infrastructure without the need for custom integrations.

ElizaOS support enables autonomous agents with Hedera capabilities to run as persistent background services, while MCP server integration brings blockchain functionality directly into development environments, such as Cursor and Claude Desktop.

Supported models

OpenAI Claude Mistral AI LLaMA by Meta

Developer experience is optimized for rapid onboarding through a modular architecture. The **plugin system separates functionality into reusable components** that instantly work across all supported. The design is already demonstrated through plugins for popular decentralized finance applications on Hedera, such as decentralized exchanges SaucerSwap and MemeJob, and decentralized lending protocol Bonzo Finance, enabling out-of-the-box functionality for automated workflows within Hedera's DeFi ecosystem.

Plugins contain reusable business logic, while adapters translate plugins into framework-specific tools.

The official @hashgraph/sdk handles actual network interaction, transaction signing, and validation. This layered approach enables developers to work at their preferred level of abstraction; whether through natural language for simplicity, direct API access for control, or custom plugins for advanced domain-specific functionality.

HIP-991: protocol-level revenue infrastructure

HIP-991 (Hiero Improvement Proposal) transforms Hedera Consensus Service from a simple messaging layer into an autonomous revenue engine for AI agents and developers. It introduces optional fixed fees for topic message submissions, enabling native monetization without relying on complex smart contracts or external payment systems.

With HIP-991, AI agents can operate as economically self-sufficient entities: collecting revenue using topics with custom fees, monitoring incoming payments, processing requests, and receiving compensation automatically, all without human intervention.

The business potential spans a wide range of applications: data marketplaces monetizing real-time information streams, AI-as-a-Service platforms where developers deploy agents charging per-query, DAO governance systems that require proposal fees, and enterprise communication networks with native billing.

Hedera's architecture stands out for its fixed predictable fees versus volatile gas costs, real-time settlement, and protocol-level integration that eliminates smart contract attack surfaces. This native monetization framework removes a key barrier preventing AI agents from acting as independent economic actors, transforming them from cost centers into autonomous, revenue-generating services.

Getting started Hedera Agent Kit

Decentralized infrastructure for sovereign AI

The partnership between Hedera Foundation, Accenture, and EQTY Lab delivers verifiable AI governance infrastructure purpose-built for public sector deployment. It addresses the heightened demand for **oversight**, **accountability**, and **reliability** as public institutions scale their adoption of AI and agentic solutions. Accenture is developing implementation blueprints and pricing models that make verifiable governance modules built on the Hedera network practical for integration across public sector ERP and digital platforms, with initial pilots planned for select AI and machine learning applications.

> The integration of these technologies not only enhances runtime governance, but also positions the Hedera network as a crucial platform. It will enable public service organizations to scale AI responsibly, leverage cryptographic solutions to reduce costs & bolster the security of deploying agentic solutions at an enterprise level.

Bryan Rich

Senior Managing Director and Global AI Lead for Health, Public Sector, and Defense

The technical architecture combines EQTY Lab's Verifiable Compute with Hedera's decentralized infrastructure to provide transparent, tamper-proof oversight of AI decision-making. This cryptographic framework enhances security and reduces deployment costs, delivering the verifiable audit trails required for high-stakes environments where transparency and integrity are mandatory, rather than optional.

The strategic integration addresses a core sovereignty challenge facing AI adoption in the public sector.
As these organizations operate across distributed digital systems, they require authority over data and decision-making processes, without dependence on centralized commercial providers. Hedera's decentralized infrastructure eliminates single points of control and enforces governance through transparent, consensus-based protocols. The result is an AI governance model that is transparent, decentralized, and enforceable by design, empowering regulated institutions to scale AI responsibly within defined legal, ethical, and geographical boundaries, and ensuring that technology ultimately serves public interests rather than consolidating power in private entities. In October 2025 at NVIDIA GTC, Hedera Foundation, Accenture, and EQTY Lab announced the launch of Verifiable Governance and Sovereignty for AI agents - integrating Hedera's enterprise blockchain infrastructure with EQTY Lab's AI Guardian solution and Accenture's multi-agent gateway to enable transparent, accountable governance in autonomous AI systems.

Running on NVIDIA DGX Cloud, the system provides immutable verifiability of nation-state and regional sovereignty controls for AI deployments across government, public safety, and critical infrastructure.

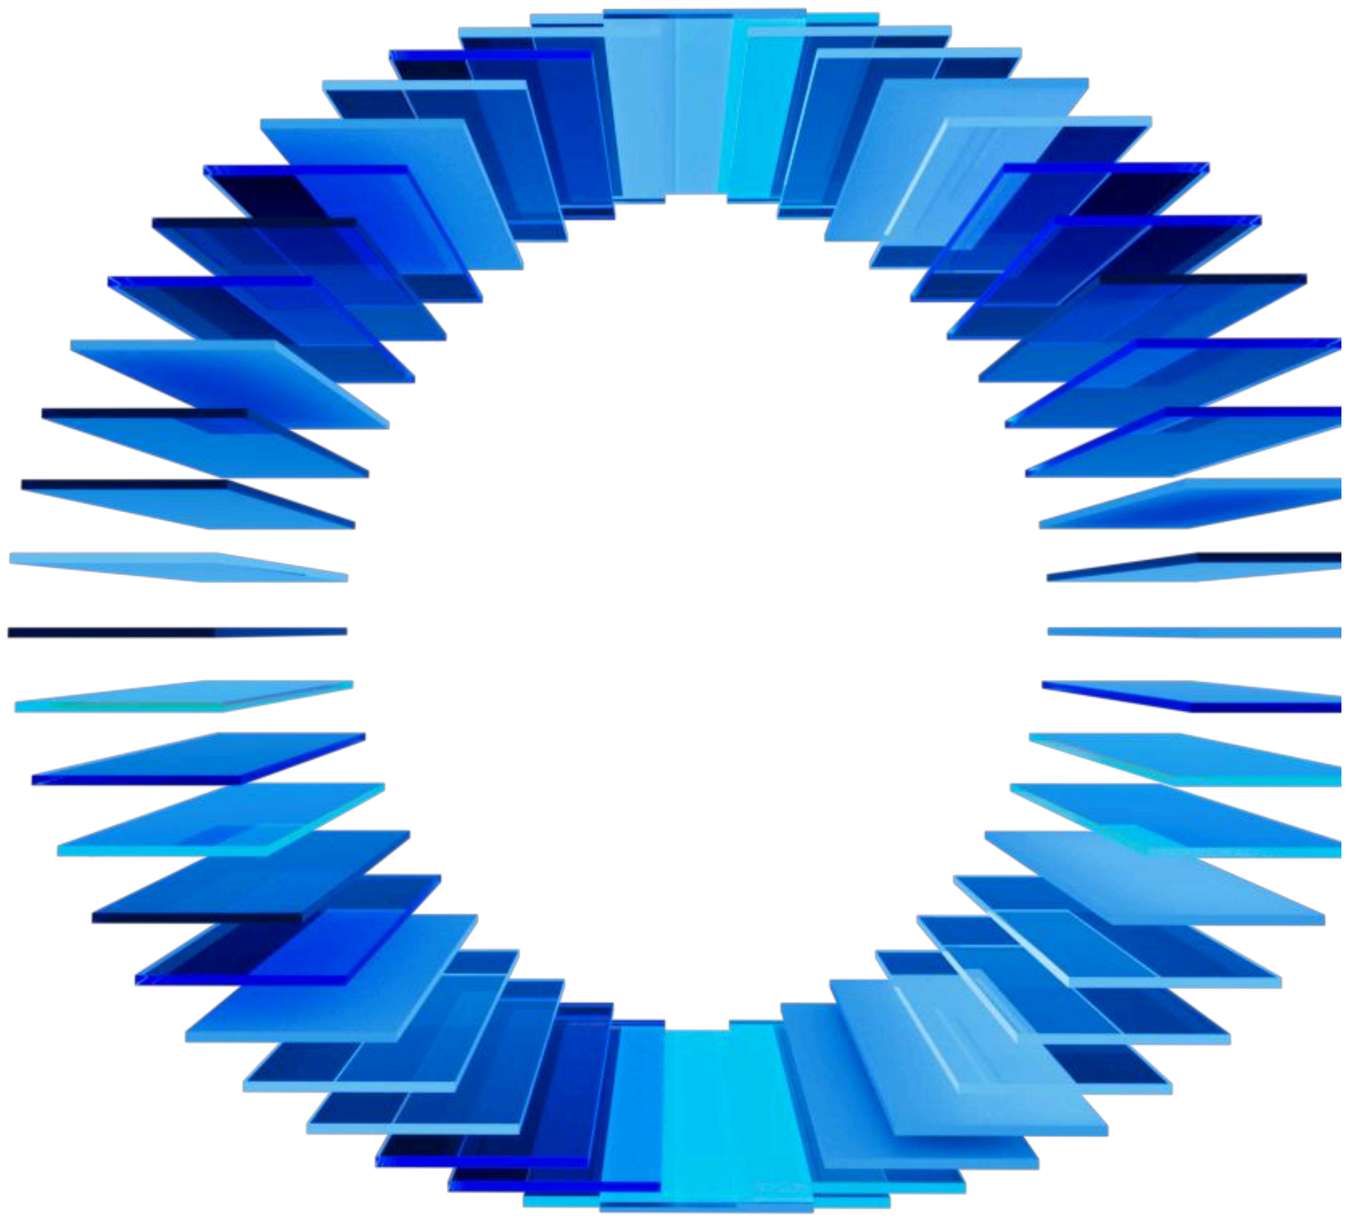

The solution introduces a digital trust layer that anchors AI agent orchestration onto Hedera Consensus Service, creating tamper-proof attestations for all agent actions and decisions. This enables government agencies and regulated enterprises to verify compliance with governance policies at runtime through encrypted attestations. Each AI agent is registered on Hedera's blockchain with verifiable records of compliance, provenance, and lineage maintained throughout its lifecycle.

Why AI on Hedera

The convergence of AI and blockchain is reshaping the digital economy. As autonomous agents evolve from simple assistants into autonomous economic actors, they demand infrastructure that traditional systems cannot deliver. Hedera serves as the trust layer for this transformation,offering cryptographic proof of computation, immutable audit trails, native tokenization, and large-scale autonomous coordination.

Scalability

Thousands of transactions per second with rapid finality, enabling real-time agent coordination at enterprise scale.

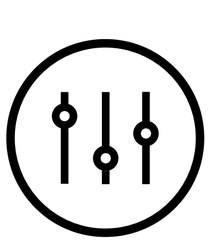

Predictability

Fixed fees denominated in USD eliminate gas volatility, allowing AI agents to model costs precisely and operate with financial confidence.

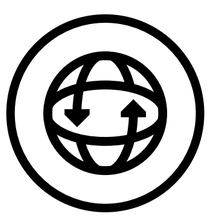

Fairness

Enterprise governance & provable fair ordering ensure transparent, accountable infrastructure without centralized control or anonymous validators.

The AI economy relies on more than powerful models. It needs trusted infrastructure where autonomous systems can transact, coordinate, and operate with cryptographic certainty. Through verifiable compute, native network services, and enterprise-grade governance, Hedera establishes the foundation for this trust.

As billions of AI agents join the digital economy, those built on transparent, scalable, and fair infrastructure will define the new standard.

AI built on Hedera is AI you can trust.

Ready to build the future?

Explore Hedera AI Studio to deploy autonomous agents in minutes, access pre-built integrations with leading AI frameworks, and join developers pioneering the agentic economy.

About Hedera

As the trust layer of the digital economy, Hedera empowers builders to create real-world impact. Hedera stands apart as the only public network governed by some of the world's most respected institutions. Built for speed, security, and scalability, Hedera provides a trusted platform for decentralized applications across highly regulated industries such as finance, supply chain, energy, healthcare, and government.

Its open-source network combines high-throughput technology, fixed low-cost fees, and real-time transaction ordering, delivering predictable performance and fairness without compromising compliance. With a thriving ecosystem and strong developer tools, Hedera is driving innovation in tokenization, AI, digital identity, DeFi, and sustainable finance.

For more information, visit **hedera.com**, or follow us on **X**, **LinkedIn**, or **Discord**. The Hedera whitepaper can be found at **hedera.com/papers**.

hedera.com **hedera** **hedera-network** **hederahashgraph**