Automata 2.0: The Machine Trust Layer for Verifiable Computation

Abstract

The proliferation of decentralized systems has exposed a fundamental tension between the need for computational integrity and the inherent transparency of public ledgers. This has given rise to systemic challenges, including value extraction, privacy violations, and a persistent trust gap that impedes the development of sophisticated, autonomous on-chain applications. Automata 2.0 introduces a modular, hardware-rooted attestation layer designed to bridge this gap. As a streamlined set of on-chain primitives, Automata 2.0 leverages Trusted Execution Environment (TEE) to establish a verifiable foundation for machine trust.[1] This whitepaper details a composable architecture of smart contracts, open-source SDKs, and APIs deployed on existing Ethereum rollups. The core of this architecture is the Data Center Attestation Primitives (DCAP) protocol,[2] which translates hardware-level cryptographic proofs into immutable on-chain attestations, making machine identity a composable primitive. We posit that this verifiable trust layer is the critical infrastructure required for the next generation of decentralized applications, with a specific focus on two high-value verticals: securing the operational integrity of autonomous AI agents and engineering fair, efficient financial infrastructure. By providing verifiable guarantees of code integrity, state confidentiality, and fair process execution, Automata 2.0 enables a new class of on-chain actors that can operate with a degree of security and autonomy previously unattainable. This architecture is coupled with a protocol economic model centered on usage-based fees and value accrual, aligning the protocol's success directly with the economic activity it enables.

1. Introduction:
The Need for Verifiable Machine Trust

1.1. The Trust Gap in Decentralized Computation

The foundational premise of blockchain technology is the establishment of trust in a trustless environment. Through cryptographic verification and decentralized consensus, these systems enable secure transactions and state transitions without reliance on a central intermediary. However, the very transparency that underpins this security model creates inherent limitations. The public nature of transaction pools (mempools) and state exposes users to adversarial strategies, most notably front-running and other forms of Maximal Extractable Value (MEV),[3] which erode market fairness and impose an invisible tax on participants. Furthermore, the complete transparency of on-chain data, while essential for verification, is fundamentally incompatible with applications that require data confidentiality.

As the decentralized ecosystem has matured, the nature of the "trust gap" has evolved. The challenge is no longer confined to concealing transaction data; it has expanded to encompass the verification of the computation itself. With the advent of complex off-chain systems, such as decentralized oracles, co-processors, and the nascent field of autonomous AI agents, a more profound question has emerged: How can an on-chain smart contract trust the output of an off-chain computational process? This requires guarantees not only of data privacy but also of execution integrity—a verifiable assurance that the correct code was executed on the correct data without tampering.

1.2. Thesis: TEEs as a Foundational Primitive for Autonomous Systems

The central thesis of this whitepaper is that Trusted Execution Environments, when anchored to a public blockchain via a robust, decentralized, and scalable on-chain attestation mechanism, transcend their role as a mere privacy tool.[4] They become a foundational and

composable primitive for building a new generation of autonomous, verifiable, and economically potent on-chain systems. TEEs furnish a hardware-based root of trust, a capability that has been the missing link for enabling smart contracts to securely interact with complex, high-performance off-chain computation and sensitive real-world data. By providing cryptographic guarantees of what code is running, on what class of hardware, and that its execution state remains confidential and untampered, TEEs empower developers to build systems that are not just decentralized, but demonstrably honest.

2. The Automata 2.0 Architecture: A Modular Attestation Layer

2.1. A Simplified Paradigm: Moving Beyond Bespoke Networks

The architecture of Automata 2.0 is defined by a strategic simplification that prioritizes capital efficiency, security, and developer composability. It avoids the operationally complex and capital-intensive model of bootstrapping a bespoke network of validator nodes. The requirement to incentivize a new, permissionless set of staking nodes introduces significant economic overhead and a "cold start" problem for security. In contrast, the 2.0 architecture leverages the mature security and infrastructure of the Ethereum ecosystem by deploying as a suite of smart contracts and services on Ethereum and Layer 2 rollups.[5] This approach allows Automata to inherit the underlying security of Ethereum's consensus, obviating the need for a separate token-based security model and allowing the protocol to focus entirely on its core logic.

The Automata 2.0 architecture consists of three primary, interconnected components:

1. **On-Chain Attestation Contracts:** A suite of smart contracts, centered around the DCAP verification logic, deployed across the Automata rollup and other supported blockchains. These contracts function as the public and immutable source of truth for machine attestation. They provide a universal, on-chain registry where any

application can query and verify the cryptographic proof of a TEE's identity and the integrity of the code it is executing.

2. **TEE-Powered Applications & Services:** A collection of both off-chain services and on-chain applications that leverage the attestation contracts to enforce specific rules and guarantees. This includes foundational infrastructure like the 1RPC verifiable relay, which provides attested access to blockchain APIs, as well as higher-level applications such as confidential order books or platforms for deploying and managing autonomous AI agents.
3. **Open-Source SDKs and Tooling:** A comprehensive set of developer libraries that abstract the complexity of TEE attestation and on-chain verification. These SDKs provide a unified interface for various TEE technologies, enabling developers to easily integrate hardware-level trust into their applications and compose TEEs with other cryptographic primitives like Zero-Knowledge Proof (ZKP) and Multi-Party Computation (MPC).

2.2. The Attestation Layer: Bridging Hardware Identity and On-Chain Logic

At the conceptual core of the Automata 2.0 architecture lies the attestation layer. This layer serves as the crucial conduit between the physical, hardware-enforced world of TEEs and the virtual, software-defined world of blockchain smart contracts. TEEs, such as Intel SGX and TDX, are capable of generating a cryptographic signature known as a "quote" through a process called remote attestation.[2] This quote is a signed statement from the CPU itself, attesting to the identity of the software running within its secure enclave and the security state of the platform. In effect, it provides a verifiable "hardware identity."

The fundamental challenge, however, is that this hardware-level proof is not natively intelligible to a blockchain's virtual machine. Automata's attestation layer, powered by the DCAP protocol, solves this problem. It provides the on-chain logic and infrastructure necessary to parse, validate, and interpret these TEE quotes within a smart contract environment.[2] This process effectively translates a machine's hardware identity into a state that can be understood and acted upon by on-chain code.

This translation turns hardware identity into a powerful and composable

on-chain primitive. It allows a smart contract to ask and receive a cryptographically certain answer to the question: "Is this off-chain service I am interacting with running the exact code I expect it to, within a genuine and secure TEE?" This capability is transformative, enabling developers to "wire" proofs of computational integrity directly into the logic of their decentralized applications, relays, and autonomous agents.

2.3. Security Model: Multi-Prover Composition (TEE + ZK + Optimistic)

Automata 2.0 adopts a sophisticated, defense-in-depth security model that acknowledges the unique strengths and weaknesses of different proof systems. The central philosophy is that reliance on any single proving mechanism creates a potential single point of failure. To mitigate this risk, Automata champions a "multi-prover" approach, a methodology for hardening Layer 2 rollups by combining multiple, independent proof systems like TEEs, ZK proofs, and optimistic fraud proofs.[6] If these heterogeneous provers produce an identical result for the same state transition, confidence in its validity increases exponentially. As one example of this philosophy in practice, the Automata Multi-Prover AVS on EigenLayer leverages Ethereum's restaked economic security to bootstrap a decentralized network of TEE-based provers that provide secondary verification services for rollups.[7] This not only enhances rollup security but also creates a global, distributed, and operational TEE compute fabric as a direct byproduct, which can then be leveraged by other application verticals.

3. Core Primitives: The Bedrock of Machine Trust

3.1. On-Chain Attestation at Scale: The DCAP Protocol

The cornerstone of the Automata 2.0 architecture is its robust and production-grade protocol for on-chain TEE attestation. This capability is primarily delivered through the Automata DCAP library, an open-source, professionally audited implementation that bridges the gap between Intel's hardware security features and the verifiable logic of smart contracts.

Intel's Data Center Attestation Primitives (DCAP) is a set of tools and libraries designed to facilitate remote attestation for SGX and Trust Domain Extensions (TDX) platforms, specifically in data center environments where direct, real-time internet access to Intel's Attestation Service (IAS) may be unavailable or undesirable.[3] This model is perfectly suited for decentralized networks, where nodes are operated by independent entities across varied infrastructure. DCAP allows for a model where the necessary cryptographic materials for verification, known as "collateral" (e.g., Provisioning Certification Key (PCK) Certificates, Certificate Revocation Lists (CRLs), and TCB Info structures), can be cached and served locally.[2]

Automata's implementation takes this concept a crucial step further by bringing the verification process entirely on-chain. The Automata DCAP library is a suite of Solidity contracts that can verify DCAP-generated quotes directly within the EVM.[8] This is made possible by a critical piece of infrastructure known as the

On-chain Provisioning Certificate Caching Service (PCCS). This service caches the attestation collateral and makes it available to smart contracts, enabling them to perform a full, trust-minimized verification of a TEE quote without relying on a centralized off-chain oracle.[9]

The protocol's maturity and scalability are not theoretical but demonstrated through extensive production deployment. It is currently live across approximately 11 networks and has processed an estimated 800,000 on-chain attestations. Its reliability is evidenced by its adoption in mission-critical infrastructure, including:

- **Flashbots' Flashtestations:** Used for the on-chain verification of Intel TDX attestations, securing components of their MEV infrastructure.[10]
- **1RPC:** A verifiable RPC relay that has served over 40 billion relays to more than 5 million monthly active users, using attestation to

prove the integrity of its relay nodes.[11]

Furthermore, the protocol is designed for the future of verifiable computation, with native paths for zkVMs. The dcap-rs library provides tools for generating and verifying attestations within Risc Zero's R0[12] and Succinct's SP1 zkVM[13] environments, allowing for the composition of TEE-based confidentiality with ZK-based public verifiability. This comprehensive, battle-tested, and forward-looking implementation of on-chain attestation forms the bedrock upon which all of Automata 2.0's applications are built.

3.2. A Unified Developer Stack: Open-Source TEE SDKs for Composable Trust

A core tenet of the Automata 2.0 strategy is the empowerment of a broad developer ecosystem through the provision of high-quality, open-source tooling. The goal is to abstract away the significant underlying complexity of TEE programming and attestation, presenting developers with a unified and accessible stack for building verifiably trusted applications. This commitment is demonstrated by a suite of maintained SDKs and libraries that provide comprehensive support for a wide range of TEE platforms, ensuring developers are not locked into a single hardware vendor. This includes robust tooling for:

- Intel SGX and TDX
- AMD SEV-SNP
- AWS Nitro Enclave
- NVIDIA H100 Confidential Compute

This hardware-agnostic approach is complemented by a philosophy of **pragmatic interoperability**. The Automata stack is not designed to be a monolithic solution but rather a set of composable building blocks that can be integrated with other privacy-enhancing technologies. The guiding principle is to allow developers to select the optimal combination of tools for their specific security, performance, and cost requirements. This composition follows a clear pattern:

- **TEE** is used to attest to the integrity of code and the confidentiality of inputs and execution state.
- **ZKP** is used to make the results of TEE computations publicly

verifiable without revealing the underlying private data.
- **MPC** is used to distribute trust across multiple, independent TEEs, eliminating single points of failure.

This composable approach is not merely theoretical but is already being realized in practice through collaborations such as with **World's AMPC**, which is exploring the use of TEEs with MPC, and the direct integration of DCAP verification into zkVMs like Risc Zero's R0 and Succinct's SP1. By providing this unified tooling, Automata allows development teams to assemble sophisticated trust architectures according to their unique risk and cost envelopes, without needing to reinvent the entire attestation and verification stack from first principles.

While TEE hardware itself is a commodity produced by major technology corporations like Intel and AMD, the software layer required to make these secure environments useful, secure, and interoperable with decentralized systems is highly specialized and complex.[14] Any project can theoretically utilize an Intel SGX chip, but building, auditing, and maintaining the intricate software stack for remote attestation, on-chain verification, and cross-platform compatibility presents a formidable technical barrier. Automata's existing suite of open-source, security-audited, and production-proven SDKs constitutes a significant strategic advantage. This pre-built and unified developer stack creates a powerful moat based on developer experience and reduced friction. As developers seek to build the next generation of trusted applications, they will naturally gravitate towards the platform that provides the most robust, accessible, and comprehensive tools, making the SDKs themselves a core product that drives ecosystem growth and adoption.

4. Application Vertical I: Trust Infrastructure for Autonomous AI Agents

4.1. The Agentic Challenge: Securing On-Chain AI

The emergence of Large Language Models (LLMs) and advanced AI has catalyzed a new paradigm: autonomous on-chain agents. These are software entities capable of controlling their own wallets, managing assets, interacting with DeFi protocols, and executing complex strategies without direct human intervention. While the potential for such agents to automate and optimize a vast range of on-chain activities is immense, it also introduces a profound and novel security challenge.

Delegating control of significant financial assets to an AI agent necessitates a level of trust that fundamentally exceeds the guarantees provided by traditional smart contracts. The agent's core logic and decision-making processes often reside off-chain, running on centralized infrastructure that is opaque to the user and the blockchain. This creates a critical "trust triangle" dilemma:

- **The User** cannot trust that the agent's code has not been maliciously altered by its developer or the infrastructure provider.
- **The Developer** cannot prove to the user that their agent will behave as advertised and that they do not have a backdoor to access user funds.
- **The Ecosystem** cannot distinguish between legitimate, well-behaved agents and malicious bots, creating systemic risk.

Without a mechanism to cryptographically enforce the integrity and confidentiality of an agent's operation, the entire field of agentic AI on-chain is relegated to experimental use cases with limited economic significance. A robust trust infrastructure is a prerequisite for their widespread adoption.

4.2. TEE-Secured Agents: A Threefold Guarantee

Automata 2.0 addresses this agentic challenge by leveraging its TEE-based attestation layer to provide a threefold set of cryptographic guarantees. This framework transforms AI agents from untrusted black boxes into verifiable and secure on-chain participants.

4.2.1. Verifiable Execution: Proof of Code Integrity

The most fundamental guarantee provided by a TEE is **remote attestation**, which serves as cryptographic proof of the software's integrity. Before a user delegates assets or permissions to an AI agent, they can demand an attestation from the TEE in which the agent is running. This attestation contains a secure hash (e.g., MRENCLAVE for SGX) of the agent's code. The user can then verify this hash against the known, open-source code of the agent. By using Automata's on-chain attestation contracts, this verification can be performed by a smart contract, which can then programmatically grant permissions only to agents that are verifiably running the correct, non-malicious code. This completely eliminates the risk of the infrastructure provider or a malicious actor surreptitiously modifying the agent's logic.

4.2.2. Confidential State and Inference: Protecting Agent Secrets

Autonomous agents must manage highly sensitive data, including private keys for their wallets, API credentials for accessing external services, proprietary trading algorithms, and private user data used for inference. TEEs provide a hardware-enforced confidential computing environment, an encrypted memory region (enclave) where both the code and the data it processes are protected from observation, even by the host system's operating system or a privileged administrator. This allows for the secure management of an agent's most critical secrets. For instance, an agent's wallet keys can be generated inside the TEE and never leave the enclave in plaintext, a model utilized by TEE-based key management solutions.[13] This ensures that even if the host server is completely compromised, the agent's assets remain secure. Similarly, confidential inference allows an agent to process private user data to generate personalized results without ever exposing that data to the underlying infrastructure.

4.2.3. Attested Access: Securing Agent Interactions via 1RPC

The trust guarantees for an agent must extend beyond its execution environment to its interactions with the outside world. An agent's

transactions and data queries are vulnerable to being monitored, censored, or manipulated by the RPC provider it uses to communicate with the blockchain. Automata's **1RPC**, a verifiable and privacy-preserving RPC relay, can be extended to serve as a secure communications gateway for AI agents. By integrating TEE attestation, 1RPC can establish a trust relationship with the agent. The agent, running in its TEE, can present an attestation to the 1RPC node, also running in a TEE, to prove its identity and integrity. This creates an end-to-end attested channel, ensuring that only verified agents can submit transactions or query data, and that their requests are shielded from eavesdropping and tampering. This model can be generalized to secure access to any external API, including LLM APIs, trading venue APIs, or MPC service APIs, creating a comprehensive security perimeter for the agent's entire operational lifecycle.

This TEE-based framework provides a direct and powerful solution to the classic "Principal-Agent Problem" as it applies to on-chain AI. In economics, this problem arises when a principal (the user) delegates authority to an agent (the AI), but cannot be certain that the agent will act in the principal's best interests. TEEs create a cryptographically enforceable "contract" between the principal and the agent. Remote attestation allows the principal to verify the exact code and rules the agent will follow before delegating control. Confidentiality and integrity guarantees ensure that neither the agent's logic nor its assets can be co-opted by a third party (the infrastructure provider). This transforms the relationship from one based on blind trust in the developer or operator to one based on verifiable, hardware-enforced guarantees, aligning incentives and making truly autonomous on-chain delegation possible.

4.3. Enabling Agentic DeFi: A New On-Chain Economy

With a foundational trust layer in place, AI agents can transition from novelties to primary economic actors, unlocking a new and sophisticated domain of decentralized finance: **Agentic DeFi**. This paradigm envisions an on-chain economy where complex financial services are managed and executed by autonomous, verifiable agents. The possibilities include:

- **Autonomous Asset Managers:** Agents that execute complex, multi-protocol yield farming or portfolio rebalancing strategies based on user-defined goals and risk parameters.
- **Intelligent Liquidity Provision:** Agents that dynamically adjust liquidity positions across multiple DEXes and lending protocols to optimize for fees and minimize impermanent loss.
- **AI-Powered DAOs:** Organizations where routine governance and treasury management tasks are delegated to verifiable agents that operate within strict, smart-contract-enforced constraints.
- **Decentralized Underwriting and Risk Assessment:** Agents that can process confidential data to make sophisticated risk assessments for insurance or lending protocols without compromising user privacy.

This emerging on-chain economy, powered by trusted agents, represents a significant new source of transactional volume and protocol revenue. The fees generated from these agentic services—for execution, for data access, for strategy deployment—will form the economic backbone of the Automata 2.0 protocol, creating a self-sustaining flywheel where the protocol's infrastructure enables a new economy, and the activity of that economy, in turn, funds and grows the protocol.

5. Application Vertical II: Fair and Efficient Financial Infrastructure

5.1. Addressing Systemic Inefficiencies: MEV and Information Asymmetry

The transparent nature of public blockchains, while a cornerstone of their security, has given rise to systemic inefficiencies in on-chain financial environments. The public mempool, where pending transactions are visible before being confirmed, creates a fertile ground for predatory strategies collectively known as MEV. Practices like front-running, where an attacker observes a user's large transaction and places their own order first to profit from the anticipated price movement, and sandwich

attacks, where an attacker brackets a user's transaction with their own buy and sell orders, are rampant.[15]

These activities are not benign arbitrage; they represent a direct extraction of value from ordinary users, creating an environment that is fundamentally unfair and economically inefficient. This information asymmetry undermines trust and acts as a significant barrier to entry for both retail and institutional participants. The objective of building fair financial infrastructure is to mitigate these issues by creating systems that can guarantee properties like **temporal fairness** (orders are processed in the order they are received) and **confidentiality** (order details are not exposed to potential attackers before execution).[16]

5.2. TEE-Powered Financial Primitives

Automata 2.0's TEE-based infrastructure provides a powerful toolkit for constructing a new generation of fair and efficient financial primitives. By moving critical parts of the transaction lifecycle into a confidential and verifiable computing environment, it is possible to eliminate the information leakage that enables MEV.

5.2.1. Confidential Order Books and Transaction Relays

The most direct way to combat front-running is to deny attackers the information they need to execute their strategies. By leveraging TEEs, it is possible to build **confidential transaction relays** and **off-chain order books**. In this model, users would encrypt their transaction or order details with a public key whose corresponding private key is held securely within a TEE enclave.

The TEE-based relay, an evolution of the 1RPC service, would receive these encrypted transactions, forming a "dark pool" or private mempool. Inside the secure enclave, the relay can decrypt the transactions and order them according to a pre-defined, fair ordering policy (e.g., first-in-first-out, or FCFS). Only after the sequence of transactions for a block is finalized is the ordered list of plaintext transactions released to a block builder for execution. Because the contents, size, and price of a

transaction are never exposed publicly while it is pending, front-runners are blinded and cannot position their own trades advantageously.

5.2.2. Verifiably Fair Order Matching Engines

Beyond just ensuring confidential submission and fair ordering, TEEs can be used to guarantee the integrity of the trade matching process itself. A decentralized exchange can run its entire order matching engine within a TEE. Through remote attestation, the exchange can provide a public, cryptographic proof to all of its users that it is running a specific, audited, and verifiably fair matching algorithm (e.g., price-time priority).[17]

This prevents any possibility of manipulation by the exchange operator, who in a traditional or opaque system could secretly alter the matching logic to favor certain participants or extract value for themselves. This approach of using TEEs to build provably fair exchanges has been explored in pioneering academic research like Tesseract[18] and is a core component of advanced industry projects like Flashbots' SUAVE, which utilizes a decentralized network of TEEs for its trust-minimized block building and MEV mitigation architecture.[19]

The application of TEEs to financial infrastructure offers a unique and pragmatic middle ground between the extremes of fully transparent and fully private systems. While cryptographic techniques like ZKPs and Fully Homomorphic Encryption (FHE) focus on proving the correctness of a final result without revealing inputs, they do not necessarily guarantee the fairness of the process that led to that result. For instance, a ZK proof can show that a set of trades was settled correctly according to an exchange's rules, but it cannot inherently prove that the initial ordering of those trades was not manipulated.

TEEs, in contrast, provide a guarantee of **Process Integrity**. Remote attestation allows participants to verify the exact algorithm being used for ordering and matching, while the confidential computing environment protects the inputs (the orders) during the critical, pre-commitment phase of the process. The final outputs (the settled trades) can then be published on-chain for public auditability. In the context of creating fair markets, the ability to cryptographically verify the integrity and confidentiality of the process is often more critical than simply verifying

the validity of the final state. This distinction positions TEEs as a uniquely powerful and well-suited technology for solving the deep-seated problems of fair ordering and matching in decentralized finance.

TEE Feature	Technical Guarantee	Problem Solved (AI Agents)	Problem Solved (Finance)
Confidentiality	Data is encrypted in-use, inaccessible to host	Protects private keys, API secrets, proprietary models	Hides order details in mempool, enables dark pools
Integrity	Code and data cannot be tampered with by host	Prevents manipulation of agent's logic or decision parameters	Guarantees the matching engine algorithm is not altered
Remote Attestation	Cryptographic proof of code hash and TEE platform state	Verifies agent is running intended, non-malicious code	Proves to all users that the exchange is running a verifiably fair matching algorithm

6. Protocol Economics and Governance

6.1. A Usage-Based Value Accrual Model

The Automata 2.0 architecture is designed around a usage-based value accrual model, often referred to as "real yield". In this model, the protocol generates revenue directly from the economic activity it facilitates. This approach aligns the protocol's success with the value it provides to the ecosystem, creating a sustainable economic engine. Revenue is derived from fees charged for the suite of services and infrastructure it provides, including:

- **Core Attestation Fees:** Fees generated from the on-chain verification of TEE attestations, the core utility of the protocol.
- **Application-Layer Revenue:** A share of fees from services built on the Automata 2.0 infrastructure. This includes revenue from

verifiable APIs, TEE-powered financial applications, and platforms for deploying and managing autonomous AI agents.

A defined portion of the protocol's net revenue, collected from these diverse sources, is then distributed to ATA token holders who actively participate in the protocol's governance. This model aligns the incentives of token holders with the long-term growth of the ecosystem and draws inspiration from the successful fee-sharing mechanisms pioneered by leading DeFi protocols. These include direct revenue distribution models, like those used by Uniswap and Curve, as well as programmatic buyback models, such as Hyperliquid, where protocol revenue is used to purchase and remove tokens from the open market.

This economic design is a direct consequence of the architectural choice to build a modular service layer that inherits its security from Ethereum. With the protocol's function being the provision of verifiable trust services that generate fees, the most logical and sustainable utility for its native token is to become the instrument for governing the protocol and capturing a share of the value it creates. This ensures the tokenomics are not an arbitrary construct but a rational component of an efficient and powerful underlying architecture.

6.2. The Role of the ATA Token

In the Automata 2.0 framework, the ATA token assumes a new set of critical functions tailored to a service-oriented, on-chain protocol:

1. **Value Accrual:** The token will serve as the direct mechanism for capturing and distributing the economic value generated by the protocol. By participating in governance (e.g., by staking or locking tokens in a voting contract), holders will be entitled to a pro-rata share of the protocol's fee revenue. This transforms the token from a security bond into a claim on the protocol's cash flow, directly tying its value to the platform's adoption and usage.
2. **Ecosystem Activation and Bootstrapping**: The protocol's treasury, funded by a portion of fee revenue and governed by token holders, can use the ATA token as a strategic tool to catalyze growth. This could involve using tokens to provide initial liquidity for a new TEE-based DEX built on the platform, offering grants to developers building innovative AI agents, or incentivizing early users to adopt

new applications within the ecosystem, thereby accelerating the flywheel of activity and fee generation.

3. **Governance:** The primary function of the ATA token will be to facilitate decentralized governance over the protocol. Token holders will have the power to vote on key protocol parameters. This includes decisions on the fee structures for various services, the percentage of revenue to be distributed versus retained in the treasury, the allocation of treasury funds for ecosystem grants and development, and the prioritization and integration of support for new TEE technologies and blockchain environments.

7. Conclusion: Building the Foundational Layer for a Machine-Driven Future

Automata 2.0 represents a significant step in the pursuit of a more secure, efficient, and trustworthy decentralized web. By embracing a modular architecture, the protocol provides a critical piece of missing infrastructure: a generalized, developer-friendly, and hardware-rooted machine trust layer. This layer, built upon the bedrock of on-chain TEE attestation, addresses the fundamental trust gap that has, until now, limited the potential of complex and autonomous on-chain systems.

The core primitives of Automata 2.0—the battle-tested DCAP protocol for on-chain verification and a unified stack of open-source TEE SDKs—transform the abstract security guarantees of trusted hardware into tangible, composable building blocks for developers. This infrastructure directly enables the next wave of decentralized innovation. For the burgeoning field of autonomous AI agents, it provides the threefold guarantee of verifiable execution, confidential state, and attested access, solving the principal-agent problem and allowing these agents to become trusted economic actors. For decentralized finance, it offers the tools to engineer verifiably fair and efficient financial systems, mitigating the value-extractive forces of MEV and information asymmetry that have plagued the ecosystem.

The protocol's economic model, based on usage-driven, fee-based value accrual, creates a sustainable and self-reinforcing system. The success of

the applications built on Automata's trust layer directly translates into value for the protocol and its stakeholders, aligning incentives across the entire ecosystem.

Ultimately, Automata 2.0 is more than a single product or protocol; it is a foundational contribution to the standards and tools that will underpin a future where computation is not only decentralized but demonstrably honest. By providing the means to verify the integrity and confidentiality of machine operations, Automata is laying the groundwork for a more secure, autonomous, and economically powerful decentralized world.

References

[1] Trusted Execution Environment - Wikipedia, https://en.wikipedia.org/wiki/Trusted_execution_environment

[2] Intel® SGX Data Center Attestation Primitives (Intel® SGX DCAP), https://www.intel.com/content/dam/develop/public/us/en/documents/intel-sgx-dcap-ecdsa-orientation.pdf

[3] Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges, https://arxiv.org/pdf/1904.05234

[4] Automata Lightpaper, https://docs.ata.network/research/lightpaper

[5] Ethereum Layer 2, https://ethereum.org/layer-2/

[6] Scaling Security: Multi-Prover Implementation on Scroll, accessed on, https://scroll.io/blog/scaling-security

[7] Multi-Prover AVS (EigenLayer) - Automata Docs, https://docs.ata.network/tee-overview/multi-prover-avs-eigenlayer

[8] Automata DCAP Attestation, https://github.com/automata-network/automata-dcap-attestation

[9] Automata Onchain PCCS, https://github.com/automata-network/automata-on-chain-pccs

[10] Flashtestations by Flashbots, https://github.com/flashbots/flashtestations

[11] 1RPC, https://1rpc.io/

[12] risc0/risc0: RISC Zero is a zero-knowledge verifiable general computing platform based on zk-STARKs and the RISC-V microarchitecture. - GitHub, https://github.com/risc0/risc0

[13] SP1 is a zero-knowledge virtual machine that proves the correct execution of programs compiled for the RISC-V architecture. - GitHub, https://github.com/succinctlabs/sp1

[14] Trusted Execution Environments (TEEs): A primer - a16z crypto, https://a16zcrypto.com/posts/article/trusted-execution-environments-tees-primer/

[15] The Potential of Self-Regulation for Front-Running Prevention on DEXes - WEIS 2023, https://weis2023.econinfosec.org/wp-content/uploads/sites/11/2023/06/weis23-heimbach.pdf

[16] Libra: Fair Order-Matching for Electronic Financial Exchanges - arXiv, https://arxiv.org/pdf/1910.00321

[17] Market Manipulation as a Security Problem - arXiv, https://arxiv.org/pdf/1903.12458

[18] Tesseract: Real-Time Cryptocurrency Exchange Using Trusted Hardware, https://eprint.iacr.org/2017/1153.pdf

[19] The Future of MEV is SUAVE, https://writings.flashbots.net/the-future-of-mev-is-suave