Kihan

FOR IMMEDIATE RELEASE

Kihan Launches Constitutional Computing to Bring Enforceable Authority to Enterprise AI Agents

New San Francisco company introduces an architecture separating AI capability from permission to act; foundational research to appear in two-part IEEE Computer feature beginning with December's special AI Governance issue

SAN FRANCISCO, Nov. 5, 2026—Kihan™, a new technology company developing Constitutional Computing, today announced its formation and the introduction of a mathematically based architecture designed to make artificial intelligence governance enforceable at the point of execution. Constitutional Computing is designed to prevent unauthorized AI agent actions before they can affect enterprise systems.

As AI agents move beyond answering questions to independently executing transactions, modifying production systems and accessing sensitive corporate information, a serious security problem is emerging. Research published by the Cloud Security Alliance in April 2026 found that 53% of surveyed organizations reported AI agents exceeding their intended permissions, while 47% experienced an AI-agent security incident in the preceding year.

Kihan was founded to address this critical problem at its source: the difference between an AI agent's ability to perform an action and its actual authority to do so. Constitutional Computing takes its name from governance, not model training: authority is written down, signed and enforced outside the AI.

Unlike conventional AI guardrails, which frequently rely on instructions to the model, Kihan's Constitutional Computing architecture establishes an independent authorization system outside the AI itself. Each consequential action is evaluated against governing authority before execution, with a cryptographically verifiable permit bound to the specific operation and its parameters. The executing tool verifies that permit before acting. On protected execution paths, no valid permit means no execution.

The architecture combines independent authority checks, execution permits bound to exact operations, and tool-side verification, with tamper-evident records of what was authorized, refused or referred for human approval.

"AI agents are becoming extraordinarily capable, but no system, however capable, should be trusted to enforce its own limits," said Jonathan Hirshon, co-founder and CEO of Kihan. "A system that can transfer $50 million, modify a production database or release confidential information shouldn't be allowed to do so simply because it has access to the necessary tools. Telling an AI to behave itself isn't security. Authorization has to be independent of the AI, verifiable and enforceable before anything consequential happens. That's why we founded Kihan. I'm privileged to work alongside a team with extraordinary experience building critical infrastructure at scale. Together, we have the expertise to deliver these safeguards at a pivotal moment for AI."

— more —

The research underlying Kihan's architecture has been accepted for publication as a two-part feature in IEEE Computer, the flagship publication of the IEEE Computer Society. The first article, "Checking LLM Outputs Before They Become State: The Problem and the Admission Gate," is scheduled for the magazine's December 2026 special issue on AI Governance, followed in January 2027 by "Checking LLM Outputs Before They Become State: Prompt Framing Across Four Models."

Kihan's first product, Inktomi™, is being developed to bring Constitutional Computing to enterprise AI environments, including agents built on models from Anthropic, OpenAI and Google, agent frameworks such as Hermes Agent, and tool ecosystems using the Model Context Protocol (MCP). Its decision engine, already operational, evaluates specific actions against signed rules, determines whether to permit, refuse or escalate a request, and preserves a verifiable decision record. Production enforcement integrations are under development, with initial enterprise evaluations targeted for December 2026.

Designed to operate on customers' own infrastructure, Inktomi complements existing identity management, cybersecurity and AI governance technologies rather than requiring their replacement. Kihan is initially targeting financial services, insurance, healthcare, legal services and other regulated industries where unauthorized AI actions can have substantial financial, operational or legal consequences.

Kihan is also exploring integration with leading enterprise technologies, including the NVIDIA OpenShell runtime and the IBM watsonx.governance platform. The Inktomi architecture is designed for deployment across modern computing environments, with compliance-oriented capabilities addressing requirements associated with GDPR, HIPAA and other regulatory frameworks.

Kihan co-founder and CEO Jonathan Hirshon is the founding architect of ATHENA, the formal foundation of the company's Constitutional Computing architecture, and author of the forthcoming two-part IEEE Computer feature. A technology industry veteran, Hirshon has nearly four decades of experience working with leading technology companies, including Apple, Sony and Silicon Graphics.

Enterprise organizations interested in Constitutional Computing and arranging a live demonstration of Inktomi can visit www.kihan.ai.

About Kihan

Kihan™, Incorporated, headquartered in San Francisco, is pioneering Constitutional Computing, a mathematically based architecture for independent authorization and enforcement of AI agent actions. Its first product, Inktomi™, is designed to ensure consequential AI actions receive explicit authorization before execution, preserving verifiable evidence of every decision. Kihan's founding principle is simple: Authority Before Commitment™. Additional information is available at www.kihan.ai.

-30-

Kihan, Inktomi and "Authority Before Commitment" are trademarks of Kihan, Incorporated. All other trademarks are the property of their respective owners.

Media Contact: Jonathan Hirshon · press@kihan.ai