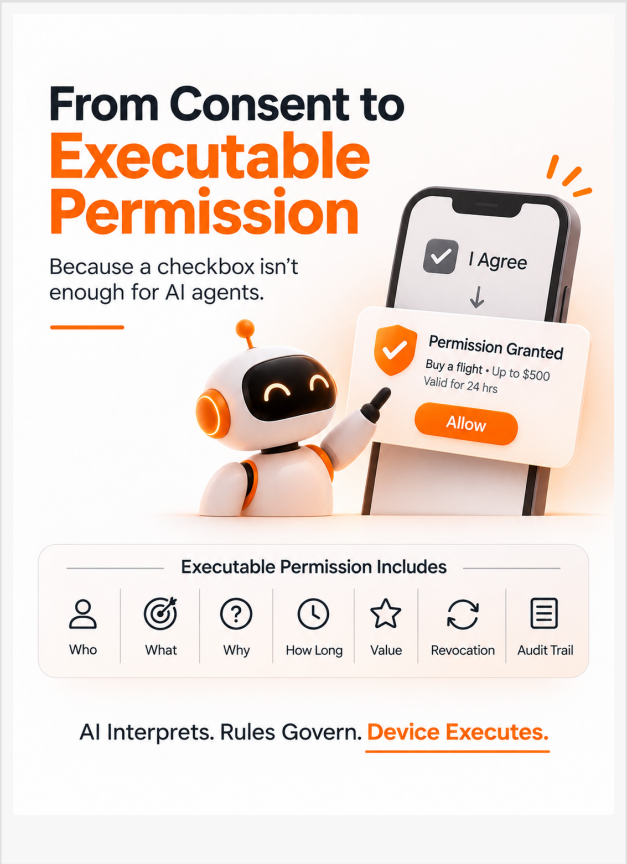
mEinstein Develops Human Context Contracts for AI Agent Authorization

AI agents need executable permission frameworks that translate human consent into defined, enforceable and revocable authority over autonomous actions.



BOSTON, MA, UNITED STATES, November 21, 2026 /EINPresswire.com/ -- As artificial intelligence agents gain the ability to access information, interact with applications and execute tasks on behalf of individuals, authorization is becoming an increasingly important consideration in AI system design. Existing consent mechanisms were largely developed for websites and digital services, while autonomous agents introduce requirements involving delegated authority, task limitations and accountability.

mEinstein is outlining a framework called Human Context Contracts to address how human authorization can be represented and applied when AI agents interact with personal information and external systems.

Traditional digital consent mechanisms generally record whether an individual has agreed to specific terms or data practices. Autonomous AI systems introduce a different set of requirements, including identifying which actions are permitted, what information may be accessed, the purpose of access and the conditions under which delegated authority remains valid.

> Permission should decay with purpose. When the approved task ends, the authority should not quietly live forever."
>
> *Prithwi R. Thakuria, Founder & CEO of mEinstein*

"A checkbox records consent," said Prithwi Thakuria, Founder and CEO of mEinstein. "An autonomous system needs permission it can actually execute."

The issue is receiving attention across the AI industry. The

National Institute of Standards and Technology (NIST) has been examining AI-agent identity and authorization, including how agents establish authority, operate under delegated permissions and maintain evidence connecting their actions to human authorization.

Microsoft has also outlined a Dynamic Permissions approach in which AI agents receive permissions appropriate to specific tasks rather than relying on unrestricted standing access. Such permissions can be limited by task and duration, with additional verification for actions requiring human approval.

OpenAI's computer-use guidance similarly distinguishes direct user instructions from information encountered by an agent during operation and recommends confirmation before certain consequential actions, including financial transactions and the transmission of sensitive information.

These developments reflect a growing distinction between recording consent and defining delegated authority that an AI system can apply during execution.

For example, a request to find a flight may authorize an AI assistant to search and compare available options. Purchasing a ticket, selecting a payment method, accepting a connecting flight or disclosing travel information to external services may require separate authorization.

Each activity involves different permissions and conditions. A single instruction does not necessarily establish authority for every subsequent action.

Thakuria describes executable permission as a structured representation of the authorized entity, permitted information or action, intended purpose, applicable limits, duration and conditions for revocation.

Commerce provides an example of how these requirements are being addressed. Mastercard's Verifiable Intent initiative focuses on establishing evidence of consumer authorization when AI agents conduct purchases. Its Agent Pay Trust Framework incorporates identity, intent, controls, execution and intelligence into the management of agent-led transactions.

Similar requirements may apply to healthcare and financial services. A healthcare system may be authorized to assess clinical-trial eligibility without automatically receiving permission to disclose patient information to a research sponsor. A financial assistant may evaluate mortgage readiness without having authority to share an individual's complete financial information or submit an application.

"Permission to know something is not permission to share it," Thakuria said. "Permission to share it is not permission to act on it."

mEinstein's proposed Human Context Contracts framework is designed to represent these

distinctions through structured requests that specify the purpose, required information and permitted action. The architecture includes authorization conditions covering the requesting entity, scope, duration, value, revocation and audit records.

The framework forms part of the company's broader Human Intelligence Infrastructure concept, which addresses the relationship between enterprise AI systems and an individual's private, user-controlled information.

Under the proposed model, an enterprise request is evaluated against the applicable permission conditions. The system can then provide an authorized result or action without automatically granting unrestricted access to the underlying personal information.

The operating principle is summarized by the company as: "AI interprets. Rules govern. Device executes."

The framework also addresses the operational limitations of repeated approval requests. Rather than requiring confirmation for every individual action, permission boundaries can define which activities an agent may perform independently and which require renewed authorization. Changes in circumstances or actions outside the approved scope may trigger additional verification.

As AI agents take on more responsibilities across digital services and commercial transactions, the distinction between consent, delegated authority and executable permission is becoming an important consideration for system developers and organizations.

Human Context Contracts represent mEinstein's approach to structuring that authorization layer, with the aim of maintaining defined human authority as autonomous systems perform increasingly complex tasks.

Krati Vyas - Global Market Strategy Leader
mEinstein
Krati.vyas@meinstein.ai
Visit us on social media:
LinkedIn
Instagram
Facebook
YouTube
TikTok
X

This press release can be viewed online at: https://www.einpresswire.com/article/957646436

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something

we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.
© 1995-2026 Newsmatics Inc. All Right Reserved.