Trust3 AI at The AI Conference 2026: AI ROI Starts with Agent Security

Trust3 AI brings agent security you can relax on to The AI Conference, Innovation Hub Booth 6.

SAN FRANCISCO, CA, UNITED STATES, September 29, 2026 /EINPresswire.com/ --
- Trust3 AI is at The AI Conference, Booth 6 in the Innovation Hub, September 30 and October 1
- The enterprise AI question is shifting to AI agent security
- Recent incidents show controls that govern access but stop short of the agent's actions
- Gartner expects more than 40% of agentic AI projects canceled by end of 2027; security review is where AI ROI stalls
- Agent DOS gives security teams the basis to approve agents: discover, observe, secure, enforced through to the data

> If an organization cannot say, in specific and auditable terms, what an agent is permitted to do, it has no basis for putting that agent into production."
>
> *Neeraj Sabharwal*

Trust3 AI will be at The AI Conference this week at Pier 48, Booth 6 in the Innovation Hub, as the enterprise conversation about AI centers around AI and agent security. California's new executive order puts frontier model safety on the public agenda. The recent AI security incidents point to the agents in enterprise production as well as to the security controls.

THE ENTERPRISE AI QUESTION HAS CHANGED
The AI question in boardrooms shifted from capability to authorization/security. An AI agent reads a record and takes an action often without a person in the loop at the moment of execution. It initiates transactions, modifies records, calls tools, and hands work to other agents. Whether an enterprise has fine-grained control over agents and actions separates a pilot from a production deployment.

RECENT INCIDENTS OF AGENTS EXCEEDING THEIR SCOPE
On September 18, Google confirmed that Gemini had accessed the systems of three real companies during a May security test. The test environment had unintended internet access, and a fictional target shared its name with a real company. The model guessed its way into one system and used credentials found in public repositories for the other two. The same day, Governor Newsom signed an executive order that proposes adding loss-of-control events to California's definition of a critical safety incident.
The order names the Hugging Face intrusion from July. An agent in a capability evaluation

escaped its sandbox, harvested credentials, and got into Hugging Face's production infrastructure. It wasn't detected until several days later.

The pattern isn't new. Back in 2025, researchers disclosed ForcedLeak, an Agentforce vulnerability in which a hidden instruction in a Salesforce lead form steered the agent into pulling CRM records and sending them to an attacker-controlled domain. The agent exfiltrated data under approved access.

In all three cases, controls governed access and stopped there. Nothing followed the agents' further actions. That is what loss of control looks like inside an enterprise.

"Agents have tremendous potential to automate many business workflows but pose enormous risk inside an enterprise if not designed in the right way. Every agent needs an identity, a declared purpose, and a policy that follows it to the data." Balaji Ganesan, Co-founder and CEO, Trust3 AI

WHY AI ROI STALLS AT THE SECURITY REVIEW

Gartner projects that more than 40% of agentic AI projects will be canceled by the end of 2027, driven in part by inadequate risk controls. Security and risk teams are blocking production because they cannot approve agents without clear visibility into them. That holds AI ROI at the pilot stage.

"Security and risk teams are not being obstructionist when they block these rollouts. They cannot approve what they cannot observe, and they cannot observe what was never instrumented. If an organization cannot say, in specific and auditable terms, what an agent is permitted to do, it has no basis for putting that agent into production." Neeraj Sabharwal, Co-founder, Trust3 AI

WHAT IT TAKES TO SAY YES TO AN AGENT IN PRODUCTION

Trust3 AI's control plane, Agent DOS, gives security teams the basis to approve agents rather than stall them:

- Discover and assess every agent for security risks, including the ones built and connected to data without a security review.
- Observe every decision, from prompt to tool call to where the output lands, so a reviewer can see what an agent did and why it was allowed.
- Secure every action with purpose-based access control (PBAC), evaluated per request, so an agent gets the access its task requires and nothing beyond it. A kill switch sits on every agent.

Because agents ultimately act on data, enforcement has to hold there too. Trust3 AI enforces policy natively at the data source on Snowflake, Databricks, BigQuery, Iceberg, and more than 50 other platforms, carrying the originating user's identity and purpose through every hop. There is no AI without data, and no AI ROI without security that follows the agent all the way through.

WHAT CUSTOMERS ARE SEEING

Here are some examples of our customers' experiences:

A Fortune 500 financial firm discovered more than 300 agents during an audit, became audit-ready without pausing operations, and now carries a continuous Trust Score on every agent.

A Fortune 100 fintech hit AWS IAM limits and replaced its over-privileged roles with attribute-

based access.
A top-5 US healthcare retailer governs thousands of BigQuery tables across hundreds of GCP projects with no ticketing queue.
Across the customer base, Trust3 AI moves teams from proof of concept to production 10x faster and cuts audit preparation time by 84%. Gartner expects that by 2029 most successful attacks on AI agents will come through access control weaknesses. The enterprises putting these controls in place now are the ones that will still be running agents then.

AT THE AI CONFERENCE
Trust3 AI will demonstrate Agent DOS (discover, observe, and secure agents and attack surfaces) live at Booth 6 in the Innovation Hub on September 30 and October 1.
Four questions worth asking before you leave the show floor:
How many AI agents are running in your environment, and what is the current security posture of each agent ?
What should a developer do to mitigate security and governance risks in their agents?
What is each agent allowed to do, what data can it access and where is that enforced?
If an agent exceeds its scope, what stops it, and how fast?
Let's chat further here: https://trust3.ai/demo/

ABOUT TRUST3 AI
Trust3 AI is the security control plane for enterprise AI agents and data. Built by the creators of Apache Ranger, the platform discovers every agent, observes every decision, secures every action, and governs every data source across any framework and any cloud.
Trust3 AI is a Snowflake Startup Accelerator member, Databricks Built On Partner, AWS Qualified Software, Microsoft Partner, Google Cloud Select Technology Partner, and NVIDIA Inception member. Investors include Point72 Ventures, Sapphire Ventures, Insight Partners, Battery, and Accel.
Learn more at https://trust3.ai.

Angela Chen
Trust3 AI
angela@trust3.ai
Visit us on social media:
LinkedIn
Instagram
Facebook
YouTube
X

This press release can be viewed online at: https://www.einpresswire.com/article/945271360

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.