ZeroTrusted.ai Launches MCP Gateway to Address Critical AI Security Risks Identified by NIST and MIT Frameworks

KISSIMMEE, FL, UNITED STATES, June 5, 2025 /EINPresswire.com/ -- ZeroTrusted.ai, the leading AI security and governance platform for enterprise environments, today announced the launch of its groundbreaking Model Context Protocol (MCP) and Agent-to-Agent (A2A) Gateway. This innovative security solution directly addresses the most critical AI risks and vulnerabilities identified in the NIST AI Risk Management Framework (AI RMF) and MIT AI Risks Repository, providing enterprises with unprecedented protection against emerging AI threats with advanced compliance intelligence and risk mitigation capabilities.

Addressing the Most Critical AI Security Challenge

The new gateway specifically targets one of the most severe risks outlined in both frameworks: AI Model Context Poisoning and Data Integrity Violations. This vulnerability occurs when malicious actors manipulate the contextual information that AI agents use to make decisions, potentially leading to compromised outputs, data breaches, and compliance violations across regulated industries.

"AI Model Context Poisoning represents one of the most insidious threats facing enterprise AI deployments today. Unlike traditional cyber attacks, these vulnerabilities can compromise the very foundation of AI decision-making while remaining virtually undetectable through conventional security measures. Our MCP Gateway provides the contextual intelligence and compliance guardrails that enterprises desperately need to protect their AI investments and maintain regulatory compliance."

— Femi Fashakin, Chief Technology Officer, ZeroTrusted.ai

Comprehensive Risk Framework Alignment

The ZeroTrusted.ai MCP Gateway is purpose-built to address the comprehensive risk categories outlined in both the NIST AI RMF and MIT AI Risks Repository:

The ZeroTrusted.ai MCP Gateway is purpose-built to address the comprehensive risk categories outlined in both the NIST AI RMF and MIT AI Risks Repository:

NIST AI RMF Compliance: Fully aligned with the Govern, Map, Measure, and Manage functions, ensuring comprehensive risk treatment and organizational oversight

MIT AI Risk Framework Integration: Direct mitigation of fairness, transparency, robustness, and privacy protection requirements

Contextual Policy Enforcement: Dynamic application of organization-specific compliance policies based on industry regulations and internal governance requirements

Ground Truth Validation: Continuous verification of AI outputs against established organizational knowledge bases and verified data sources

Beyond Traditional Gateway Limitations

While competitors rush to market with conventional security approaches that merely adapt existing gateway technologies, ZeroTrusted.ai's solution addresses the unique vulnerabilities inherent in Model Context Protocol communications and Agent-to-Agent interactions. Traditional gateways fail to understand the contextual nature of AI communications, missing critical threats that exploit the semantic relationships between AI agents and their data sources.

The platform's advanced contextual intelligence system enables security teams to define custom risk parameters that are specific to:

Unique MCP tool configurations and agent behaviors within the organization

Proprietary intellectual property and sensitive data patterns

Industry-specific compliance obligations (HIPAA, GDPR, SOX, etc.)

Organizational risk tolerance and governance policies

Contextual Intelligence for Enhanced Accuracy

The gateway's sophisticated risk assessment capabilities leverage the organization's own contextual data and established ground truths to evaluate the accuracy and appropriateness of information transmitted between AI agents. This approach ensures that security decisions are made with full understanding of the organization's specific data landscape, operational context, and risk profile.

"This contextual awareness is what sets our solution apart," explains Fashakin. "We're not just filtering traffic—we're providing intelligent governance that understands your organization's unique AI ecosystem and can prevent risks that are specific to your data, your agents, and your compliance requirements."

Enterprise-Grade AI Security Platform

ZeroTrusted.ai is an AI security and governance platform purpose-built for enterprises operating large-scale, distributed, and regulated AI systems. The platform offers real-time protection, continuous observability, and deep compliance intelligence for AI workloads running on Azure Kubernetes Service (AKS), Azure Arc, private cloud, or on-premises environments.

Built for enterprises deploying AI across cloud, on-prem, and hybrid environments, ZeroTrusted.ai provides comprehensive platform protection for AI models, applications, and agents from modern threats—including data leakage, adversarial attacks, hallucinations, and compliance violations.

About ZeroTrusted.ai

Founded by experts in AI, cybersecurity, and cloud infrastructure, ZeroTrusted.ai helps organizations enforce trust, transparency, and accountability across the entire AI lifecycle. From real-time protection and model observability to compliance automation and secure multi-agent orchestration, ZeroTrusted.ai delivers the guardrails that enterprises need to deploy AI safely and responsibly.

Trusted by global companies in finance, healthcare, logistics, and government, ZeroTrusted.ai is redefining how organizations secure the future of AI.

Availability and Contact Information

The ZeroTrusted.ai MCP Gateway is available immediately for enterprise customers. For more information about ZeroTrusted.ai and its comprehensive AI security platform, visit www.zerotrusted.ai.

Sharon Lam
ZeroTrusted.ai
+1 407-507-9350
email us here
Visit us on social media:
LinkedIn
Instagram
Facebook
X

---

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.