Salt Security Launches Industry's First Agentic Security Platform for the AI Stack Across LLMs, MCP Servers and APIs

PALO ALTO, CA, UNITED STATES, March 18, 2026 /EINPresswire.com/ -- New platform gives enterprises full visibility and governance across the Agentic Security Graph connecting LLMs, MCP servers and APIs, helping security teams understand not just what AI agents can say, but what they can do

Salt Security, the leader in API security, today announced the launch of the Salt Agentic Security Platform, a new platform designed to secure the rapidly expanding Agentic Security Graph inside modern enterprises and enable organisations to adopt AI agents safely and at scale.

As enterprises deploy AI agents to drive greater efficiency and productivity, their success depends on more than model quality alone. It depends on how effectively those agents connect to enterprise systems, data and workflows. The more connected an agent is, the more valuable it becomes, and the more security risk it can introduce. To understand that risk, security teams need visibility into the full set of relationships between LLMs, MCP servers, and APIs that enable agent behaviour.

A simple way to think about an AI agent is as a digital employee:

LLM: The brain - responsible for reasoning and decision-making.
MCP servers: The hands - enables the agent to press buttons and pull levers
APIs: The buttons and levers - allow the hands to take real action across enterprise environments.

Together, LLMs, MCP servers and APIs form the pillars of what Salt calls the agentic stack. Mapped together, they create the Agentic Security Graph, a security context layer that helps organisations understand how agents reason, connect, and take action across the enterprise.

APIs are the action layer that allows AI agents to interact with enterprise systems, execute workflows, and access sensitive data. When those interactions are not properly secured, an AI agent can move beyond generating responses to taking actions that can cause real operational harm.

As organisations accelerate adoption, the number of LLM connections, MCP servers, APIs and

autonomous interactions across enterprise environments is growing exponentially. The Salt Agentic Security Platform gives security teams a unified way to discover, visualise, govern and protect this entire Agentic Security Graph rather than just individual components within it.

“Most AI security solutions focus on prompts and models,” said Roey Eliyahu, CEO and co-founder of Salt Security. “But the real enterprise risk is not just in what an agent can say. It is in what an agent can do through MCP servers and APIs. These systems connect agents to data, workflows and enterprise services. That is what we call the Agentic Security Graph, and Salt’s Agentic Security Platform is designed to expose and secure it.”

To secure the agentic fabric, Salt introduces two new security capabilities:

Agentic Security Posture Management (AG-SPM) and Agentic Detection and Response (AG-DR), creating a unified approach to securing the full agentic lifecycle from code to runtime.
Agentic Security Posture Management (AG-SPM): Continuous discovery and governance of LLM connectivity, agents, MCP servers, APIs and the relationships between them.
Agentic Detection and Response (AG-DR): Real-time detection of abuse, misuse, and anomalous behaviour across LLM connectivity, agent-driven activity, MCP servers and APIs.

Together, these capabilities help organisations adopt AI agents safely, turning security from a potential blocker into a true enabler of business innovation.

Customer Validation
Early customers are already benefiting from using the platform to gain visibility into rapidly expanding AI environments:

“As we deploy more AI agents across our organisation, the complexity of the systems they interact with has increased dramatically and is challenging to manage,” said Shawn Griffin, Chief Information Security Officer, CFIUS Security Officer & Cybersecurity Officer at Siemens. “Salt is uniquely positioned to secure this new environment because every agent interaction ultimately runs through APIs. The Agentic Security Platform has already given us improved visibility and protection that we need to confidently scale AI across the Siemens Software business.”

Map Your Agentic Security Graph
Salt Security is offering a limited number of Agentic Security Graph Discovery Sessions during the RSA Conference that reveal how LLMs, MCP servers and APIs connect across your enterprise. In minutes, Salt reveals the hidden API relationships powering your AI systems.

Availability is limited. Request your discovery session here.

For more information, visit www.salt.security.

About Salt Security

Salt Security developed the industry’s first Agentic Security Platform, designed to secure the agentic infrastructure created by LLMs, AI agents, MCP servers and APIs. As organisations deploy autonomous agents that interact with enterprise systems through APIs, traditional security boundaries break down. The Salt platform continuously discovers, governs and protects this new architecture, providing visibility into agent capabilities, API interactions and sensitive data exposure across the agentic enterprise from code to runtime.

Bethany Smith
Eskenzi PR
beth@eskenzipr.com
