

SOLUTION BRIEF

Securing the Full Agentic Stack with the Salt Agentic Security Platform

The Challenge: Visibility Gaps in the Agentic Stack

As enterprises deploy AI agents to drive productivity, their success depends on more than model quality alone. It depends on how effectively those agents connect to enterprise systems, data, and workflows. This connectivity creates a new architectural layer: the agentic stack. These agents function as digital employees, using Large Language Models (LLMs) for reasoning, Model Context Protocol (MCP) servers for connectivity, and APIs for execution.

Because these agents operate at machine speed and can improvise their own workflows, traditional security tools are blind to their activity. Organizations are left with a massive visibility gap, unable to see how agents connect to internal systems or what sensitive data they access within the Agentic Action Layer.

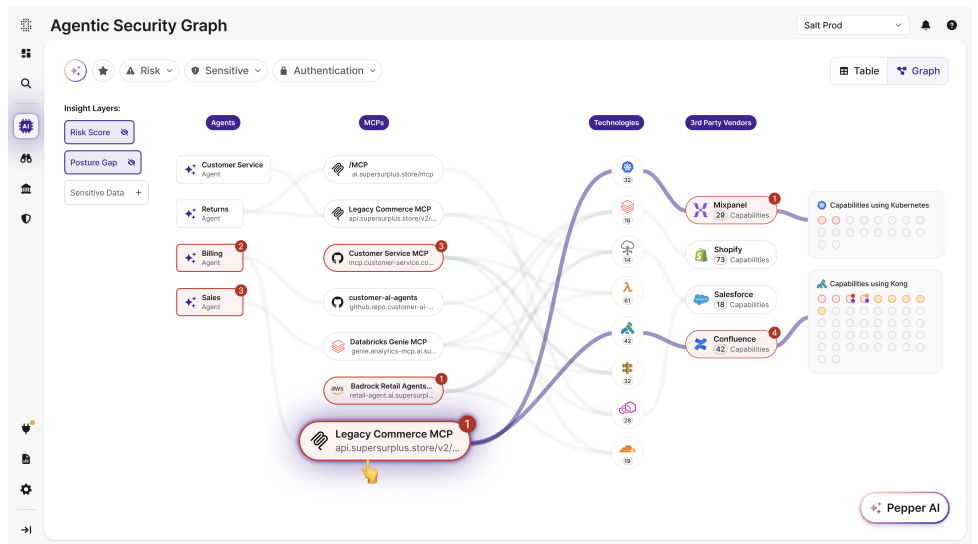
The Salt Solution: The Industry's First Agentic Security Platform

The Salt Agentic Security Platform is a dedicated solution designed to secure the interactions between AI agents and enterprise data. It provides a unified way to discover, visualize, and protect the infrastructure that powers agent behavior: the LLM connections, MCP servers, and APIs that form the **Agentic Action Layer**.

As the industry's first centralized system of record for agentic infrastructure, Salt consolidates intelligence from across the environment, capturing external exposure, analyzing code repositories, and monitoring live runtime traffic. This comprehensive approach ensures that security leaders have the context they need to govern AI with confidence.

The Agentic Security Graph: A Unified Context Layer

To solve the visibility crisis, the Salt platform provides the **Agentic Security Graph**. This is a purpose-built security context layer that serves as the central hub for visualizing every relationship discovered and monitored by the platform. It maps the complex, multi-layered web of interactions between LLMs (the brain), MCP servers (the hands), and APIs (the buttons and levers).



The Agentic Security Graph serves as the "source of truth" for unifying governance insights from AG-SPM and real-time behavioral detections from AG-DR. By delivering this integrated view, Salt allows security teams to move beyond individual logs and see exactly how an agent reasons, connects, and takes action across the enterprise.

Core Platform Capabilities

Agentic Security Posture Management (AG-SPM)

AG-SPM provides continuous discovery and governance of the agentic lifecycle from code to runtime. Using the Agentic Security Graph, Salt provides visibility into the full set of relationships that power agent behavior, ensuring every machine identity adheres to established security standards.

- **Continuous Multi-Pronged Discovery:** Automatically identify all LLM connectivity, AI agents, MCP servers, and APIs. Salt identifies forgotten or "zombie" MCPs exposed to the public internet and scans private repositories to find shadow integrations before they are deployed.
- **Regulatory Guardrails:** Establish and enforce governance policies aligned with emerging AI standards, such as the EU AI Act and ISO 42001, to ensure autonomous interactions are traceable and auditable.
- **Fabric Governance:** Eliminate "Shadow MCP" servers and ensure every agent operates within the logical boundaries of its intended business function.

Agentic Detection and Response (AG-DR)

AG-DR delivers real-time detection of abuse, misuse, and anomalous behavior across the agentic lifecycle. Since agent behavior is dynamic, Salt moves beyond static signatures to identify malicious intent at the Agentic Action Layer.

- **Agentic-Aware Baselines:** Establish behavioral norms for LLM connectivity and agent-driven activity to detect anomalous patterns, such as mass data pulls or unauthorized tool usage.
- **Identity-Aware Intent Analysis:** Correlate 100% of traffic back to the unique agentic identity to recognize when reasoning turns into unauthorized exfiltration. Salt identifies the "Sequence of Intent" to catch logic-based attacks that packet inspection misses.
- **Automated Mitigation:** Interrupt machine-speed attacks by providing real-time blocking triggers for the attacker's session, preventing material operational harm and data loss in seconds.

Why Salt Security?

Salt Security is the leader in securing the Agentic Action Layer. While model-centric tools focus on prompt filtering, the Salt Agentic Security Platform secures the infrastructure where actions are taken. By providing the **Agentic Security Graph**, Salt delivers the visibility and compliance evidence needed to safely adopt AI agents at scale.

Conclusion

The transition to an agentic enterprise represents a fundamental shift in how software executes and interacts with data. This shift requires a corresponding evolution in security visibility. The Salt Agentic Security Platform turns the opaque risks of autonomous agents into a managed, visible, and secure environment. By securing the infrastructure powering the next wave of AI innovation, Salt enables organizations to deploy agentic architectures with the confidence that their most critical systems and data remain protected within the Agentic Action Layer.

