

Salt Agentic Security Platform

See every agent. Stop every attack path. From the first prompt to the API behind it.

CATEGORY

Agentic Security

PRODUCT

AG-SPM and AG-DR with AI-DR

YEAR

2026

OVERVIEW

An attack on a model is now an attack on the business

AI agents connect language models to MCP servers, tools, and the APIs that run the business. A malicious prompt no longer stops at a bad answer. It can end in unauthorized access, a leaked record, or a refund that should never have been paid.

The Salt Agentic Security Platform gives security teams one connected view of that entire path, and the controls to stop an attack at any point along it. Salt now adds native AI Detection and Response (AI-DR) inside AG-DR, closing the last major runtime gap at the LLM layer.

49.8%

of organizations confirmed or suspect an AI agent took an action they did not intend, expect, or authorize in the past year.

Salt Security, 2H 2026 State of Agentic AI and API Security

12.5%

can consistently trace an agent's full path from the initial prompt through the MCP servers and APIs it reached.

Salt Security, 2H 2026 State of Agentic AI and API Security

#1

Prompt injection is the top ranked risk in the OWASP Top 10 for LLM Applications, in both direct and indirect forms.

OWASP Top 10 for LLM Applications, 2026

THE CHALLENGE AND THE APPROACH

The attack is one story. Security sees five.

Enterprises run LLM guardrails, AI gateways, MCP security, API security, endpoint and SASE controls, and code scanning. Each protects its own layer and none sees the whole path.

The same incident reads as prompt injection to the model layer, abusive tool use to the MCP layer, and an authorization failure to the API layer. To the business, it is one breach.

One platform that maps the path and stops the attack

Two capabilities work from a single Agentic Security Graph. AG-SPM shows teams what agents exist, what they can reach, and where the gaps are. AG-DR detects and stops attacks across those agents and every system they touch.

Existing guardrails and gateways stay in place. Salt connects them and fills the gaps they leave.

Prompt injection → compromised agent → MCP tool → backend API → business impact. One attack path, not five disconnected alerts.

PLATFORM STRUCTURE

Two capabilities, one Agentic Security Graph

AG-SPM answers what you have and what it can reach. AG-DR answers what is happening right now and stops it. Both read from the same graph, so posture informs detection and detection feeds back into posture.

AG-SPM Agentic Security Posture Management	AG-DR Agentic Detection and Response
01 Continuous discovery Finds agents, LLMs, and MCP servers from code to runtime. Shadow AI and shadow MCP surface early.	01 Native AI-DR for the LLM layer Real-time protection against direct and indirect prompt injection, jailbreaks, unsafe model behavior, sensitive data moving through LLM calls, and unsanctioned model access.
02 Agentic Security Graph Maps how agents connect to models, MCP servers, tools, downstream APIs, and data.	02 Agentic baselines Learns normal agent and LLM connectivity behavior, then flags mass data pulls and unauthorized tool use.
03 Permission and boundary control Catches over permissioned agents and keeps each one inside its intended business function.	03 Identity-aware intent analysis Ties all traffic to a unique agentic identity to tell legitimate agents from malicious bots and to catch logic based abuse.
04 One view of guardrails Brings cloud, gateway, endpoint, and SASE guardrail configurations into the graph and shows where protection is missing.	04 Attack path correlation Connects an attack on the model to the MCP and API attacks that follow, as one incident.
05 Governance and audit Keeps every autonomous interaction traceable against frameworks such as the EU AI Act, NIST, and SOC 2.	05 Automated blocking Interrupts machine speed attacks in real time, even when the attacker skips the agent and hits an API directly.

Four sources feed the graph

Salt sees an agentic environment the way an attacker does and the way the platform team does, then confirms it against live traffic.

Surface
External scan of what is exposed to the internet


Connect
Cloud, gateway, and AI platform configuration

Code
Repositories and AI coding assistants

Collect
Live runtime traffic across 70+ technologies

SEE IT IN ACTION

A billing agent attack, seen as one incident



WHY SALT IS DIFFERENT

Start at the action layer, where agents do damage

Most AI security tools start at the model and work outward. Most API tools start at the edge and ignore agents. Salt starts where agents act, calling APIs and moving data, and extends to the model. That starting point comes with eight years of runtime API security depth.

DIMENSION	TYPICAL APPROACH	SALT
Where it starts	The prompt and model layer, or the network edge	The action layer: the APIs, MCP servers, and data agents reach, then up to the LLM
What it sees	Each control sees only its own layer	One graph from prompt to model, agent, MCP server, tool, API, and data
How it discovers	From inside one agent platform, or from edge traffic only	Four sources: Surface, Connect, Code, and Collect
Existing guardrails	Replace them, or leave gaps between them	Keep them. Salt brings them into one view and adds native AI-DR where coverage is missing
How it detects	Signatures, rules, and rate limits	Behavioral baselines and identity-aware intent analysis
Focus	A module inside a broader suite	An independent, dedicated agentic and API security platform

KEY BENEFITS

Security that follows the whole path

See the full attack, fast

Security teams understand how an attack started, what it targeted, and which business systems are at risk, without stitching together five tools.

Stop machine speed attacks

Detect malicious intent instead of known signatures, and block it in real time before data leaves.

Keep what you already own

Edge controls and guardrails stay. Salt unifies them and covers homegrown agents, such as those running in Kubernetes, that sit outside them.

Prove AI governance

Give auditors and the board a traceable record of what every agent can reach and what it did.

AVAILABLE NOW

Native AI-DR is part of Salt AG-DR

Available within the Salt Agentic Security Platform alongside AG-SPM. See your own Agentic Security Graph and the gaps in your guardrail coverage.

Book a demo

salt.security/demo-request