SALT

What Is Agentic AI—and Why It Matters Now

SALT

Over the past few years, AI has evolved from simply generating content to taking autonomous action. We're no longer just asking LLMs to write text—we're embedding them into agents that initiate workflows, make decisions, and operate continuously across environments.

This is **Agentic AI:** a new software paradigm where AI agents are:

- **Goal-driven** and **autonomous**
- **Memory-enabled,** using tools like vector databases
- **API-powered**, interacting with systems through live calls
- Capable of **reasoning over time** and **adjusting behavior**

Most Agentic AI systems use **LLMs for reasoning** and **MCP servers (Model Context Protocol)** to manage the agent's tools, memory, and environment context. MCP servers act as the *API brokers*, orchestrating what agents know, what they can do, and how they operate.

Agentic AI is no longer experimental. It's already being deployed in industries like finance, healthcare, logistics, and SaaS platforms.

Why Agentic AI Creates New Security Risks

Agentic AI holds tremendous potential—but it introduces a **fundamentally different risk profile**:

Autonomy introduces unpredictability: Agents make decisions you didn't explicitly code—and may cross trust or security boundaries.

Memory retains sensitive data: Agents persist data in memory stores that can include proprietary or regulated information.

APIs are the execution layer: Agents rely on APIs to access systems, retrieve data, and trigger actions—API traffic is now autonomous, not human-initiated.

These agents interact in complex ways. They can chain actions across APIs, invoke external tools, and pass data between systems. And traditional security controls—like WAFs, API gateways, and CDNs—**don't see what's really happening**.

They miss:

- Agent identities and intent
- Memory access or leakage
- Tool usage via MCP orchestration
- Sensitive data persistence
- Shadow agents or rogue MCP servers

SALT

How Salt Secures Agentic AI and MCP Traffic

Salt Security provides the industry's first agent-aware API security platform, purpose-built to monitor and protect AI agents, MCP servers, and the dynamic API fabric they operate on.

See and Profile Agent Behavior

Salt continuously identifies:

- Which agents are active
- What they're doing
- How they interact with APIs and memory

Secure MCP Traffic and Tool Invocations

Salt observes and enforces behavior at the MCP layer, ensuring agents can't abuse tools or access out-of-scope data.

Detect Memory and API Anomalies

Salt spots:

- Prompt injection attempts
- Sensitive data misuse
- Policy violations
- Chained actions that lead to privilege escalation

Deploy Instantly with Salt Cloud Connect

Get API visibility in minutes—no traffic mirroring, no agents to install. Salt connects unobtrusively to your API infrastructure (including AWS) and provides immediate insights.

Real-World Impact: Case Snapshot

Global Airline Deployment

- 3,200 AI agents
- 8,750 agent-related APIs
- 12 MCP servers, 5 shadow servers uncovered
- Salt detected:
  - Agents accessing legacy flight systems
  - PII leakage from memory context drift
  - Abuse of baggage APIs with write access
  - Misconfigured MCP server exposed externally

Salt provided full visibility, enforced memory and tool policies, and reduced agent-related security alerts by **61%.**

SALT

You Can’t Secure AI Without Securing APIs

Agentic AI is fast-moving, high-impact—and inherently unpredictable. But it’s not enough to secure the LLM or the agent code. You must secure the **API fabric** they rely on.

Salt Security gives you the visibility, control, and protection needed to adopt Agentic AI safely and at scale.

Learn more at salt.security or request a demo.