

STRATEGIC BRIEF

CISO BRIEF: Securing the Agentic Enterprise

The Situation: A Fundamental Shift in API Consumption & Risk

The rapid adoption of AI is creating a fundamental shift in how our digital infrastructure operates. The primary consumers of our APIs are no longer human developers or traditional applications, but a new generation of autonomous AI agents. This change is not incremental; it represents a complete paradigm shift in our attack surface.

According to **Gartner®**, "By 2028, 80% of organizations will see AI agents consume the majority of their APIs, rather than human developers." (How MCP and the A2A Protocols Impact API Management - Gartner)

This massive shift directly translates to a new front line for cyber risk. Gartner also projects that, "By 2028, over 50% of API security breaches will be related to AI." Every action an AI agent takes is an API call. As we deploy more agents and connect them to data via protocols like MCP (Model Context Protocol), we are exponentially increasing the number and complexity of API interactions within our environment, creating significant and unforeseen security blind spots.

Top 3 Material Risks of Ungoverned Agentic APIs

Traditional security tools like WAFs and API gateways were not designed to understand the context and behavior of autonomous agent traffic. This creates three critical, board-level risks:

1. **Pervasive Access Control Failures:** Gartner predicts that "through 2029, over 50% of successful cybersecurity attacks against AI agents will exploit **access control issues**, using direct or indirect prompt injection as an attack vector." Agents often require broad access to function, and without granular, context-aware security, they become prime targets for attackers to manipulate and gain access to sensitive systems and data.
2. **APIs as the Primary AI Attack Vector:** APIs are not just enabling AI; they are the primary way it will be attacked. As **KuppingerCole** explains in their [Leadership Compass for API Security and Management](#), "APIs are the backbone of AI," and because of this foundational role, "Most AI-related vulnerabilities, including prompt injection, data exfiltration, or model abuse, are exposed through insecure APIs." Without securing the API layer, any security measures applied



directly to AI models are insufficient.

3. **Massive API Sprawl & Shadow API Proliferation:** The rapid pace of AI development results in the constant deployment of new agents and MCP servers, leading to a vast number of unmanaged, unmonitored, and unsecured "shadow APIs." Without a dynamic and continuous discovery process, we cannot protect what we cannot see, leaving vast portions of our infrastructure exposed.

Recommended Strategic Actions

To enable secure innovation and govern this new agentic ecosystem, we must evolve our API security strategy beyond traditional approaches. We recommend a three-pronged, vendor-agnostic strategy:

1. **Achieve Total Visibility:** Implement a continuous, real-time discovery process to maintain a complete and always-current inventory of all APIs, including those consumed by AI agents and exposed via MCP servers.
2. **Establish Agent-First Governance:** Shift from a developer experience to an agentic experience framework. This requires establishing a robust posture management program to define and enforce specific security policies for machine identities, ensuring they have the least privilege necessary.
3. **Adopt Runtime Behavioral Security:** Deploy a dedicated API security solution that uses AI-powered behavioral analysis to baseline normal agent activity. This is the only way to detect and block anomalous behavior in real time, stopping sophisticated attacks that bypass traditional defenses.

Gartner, How to Adapt Your API Strategy to Succeed in the AI Era, By Shameen Pillai, Mark O'Neill, 24 February 2025

Gartner, How MCP and the A2A Protocols Impact API Management, By [Shameen Pillai](#), [Mark O'Neill](#), [Aaron Lord](#), 25 August 2025

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

This strategic brief was developed by Salt Security, an Overall Leader in the 2025 KuppingerCole Leadership Compass for API Security and Management. We help CISOs secure their most critical innovations. Get a [Checklist](#) to help you and your team with security strategies as you work on AI projects.

