



NEWS RELEASE

F5 and MuleSoft, a Salesforce Company, Collaborate to Deliver Inline Security and Governance for Agent Fabric and Agentic AI Applications

2026-09-02

Joint integration featuring F5 AI Security Platform and Agent Fabric embeds AI runtime security directly into agentic AI workflows, blocking threats like prompt injection and sensitive data exposure

SEATTLE--(BUSINESS WIRE)-- F5 (NASDAQ: FFIV), the global leader in delivering and securing every app and API, today announced a technology integration with MuleSoft, a Salesforce company, bringing F5 AI Guardrails—part of the **F5 AI Security Platform**—directly into Agent Fabric. As enterprises rapidly scale AI agents and large language model (LLM) applications, this native integration provides security and platform engineering teams with centralized policy enforcement, real-time protection against malicious prompts, and enhanced auditability without re-architecting or replacing the systems already in use.

As organizations scale agentic AI capabilities, security teams face a growing governance gap and frequently lack visibility into agentic traffic. They risk exposure from prompt injection, harmful outputs, and sensitive data leakage. The joint integration federates F5's runtime AI security into Agent Fabric's Omni Gateway, enabling enterprises to enforce unified guardrail policies across all prompts and outputs moving through their agentic workflows.

Eliminating the governance gap for agentic AI

Until now, organizations deploying Agent Fabric and seeking to leverage **F5 AI Guardrails** have faced a trade-off:

either route LLM traffic through a separate F5 inspection layer, creating additional operational complexity and fragmented telemetry, or rely solely on native gateway controls without extending F5 AI Guardrails policies directly into their MuleSoft-managed workflows.

By federating F5 AI Guardrails into Agent Fabric, enterprises gain:

- Unified governance and zero double-proxy overhead: Policy management for AI traffic lives in a single control plane. Agent Fabric's Omni Gateway routes LLM calls directly to the F5 AI Guardrails Scan API, inspecting inbound prompts and outbound completions inline before models are invoked or responses returned.
- Support for Agentforce ecosystems: Organizations can apply consistent runtime security controls across Agentforce-powered agents, Agent Fabric workflows, and custom AI applications.
- Proactive threat mitigation: The solution is designed to block prompt injection, jailbreaks, toxicity, and unauthorized topics while reducing personally identifiable information (PII) and protected data exposure at runtime.
- Data residency and sovereign control: Flexible dual-deployment topology allows self-hosted Kubernetes deployments including private VPCs, allowing sensitive prompt and completion data to remain within customer boundaries.
- Low-touch policy tuning: Security teams author and version scanners, blocklists, and sensitivity thresholds within the F5 console, which Omni Gateway picks up dynamically without requiring policy or code changes.
- SOC back-correlation and compliance auditability: Decisions carry detailed telemetry and shared scan IDs for seamless correlation in the F5 console, simplifying compliance with regulations such as the EU AI Act, GDPR, and HIPAA.

"AI agents are moving from experiments into the critical path of the enterprise," said Kunal Anand, Chief Product Officer at F5. "The biggest risk in agentic AI is that agents will move faster than enterprise security and governance models can keep up. By integrating the F5 AI Security Platform directly into Agent Fabric, we are putting protection in the path of every prompt and response, where it can operate in real time. That lets organizations move faster with AI while preserving the control, visibility, and accountability their most important business processes demand."

"We built Agent Fabric as the neutral solution to support enterprises running agents across a mix of models and platforms," said Andrew Comstock, SVP & GM at MuleSoft. "By extending Agent Fabric's existing LLM API and AI services to support F5 AI Guardrails as a first-class provider, we're making it even easier for customers to scale their agentic enterprise on the security tooling they know and trust."

Availability and Dreamforce demonstration

F5 AI Guardrails for Agent Fabric is now generally available. At Dreamforce (September 15–17, 2026, San Francisco),

organizations can engage with the Agent Fabric team to learn more about the integration and explore how to apply AI security and governance controls across agentic AI deployments.

To learn more about how F5 and MuleSoft are securing the future of agentic AI, visit <https://www.f5.com/solutions/ai-security>.

About F5

F5, Inc. (NASDAQ: FFIV) is the global leader that delivers and secures every app. Backed by three decades of expertise, F5 has built the industry's premier platform—F5 Application Delivery and Security Platform (ADSP)—to deliver and secure every app, every API, anywhere: on-premises, in the cloud, at the edge, and across hybrid, multicloud environments. F5 is committed to innovating and partnering with the world's largest and most advanced organizations to deliver fast, available, and secure digital experiences. Together, we help each other thrive and bring a better digital world to life.

For more information visit **f5.com**

Explore F5 Labs threat research at **f5.com/labs**

Follow to learn more about F5, our partners, and technologies: **Blog | LinkedIn | X | YouTube | Instagram | Facebook**

F5 is a trademark, service mark, or tradename of F5, Inc. or its affiliates in the U.S. and other countries. All other product and company names herein may be trademarks of their respective owners. The integration described herein does not create a partnership, joint venture, or agency relationship between the parties.

Source: F5, Inc.

Jenna Becker

F5

(415) 857-2864

j.becker@f5.com

Holly Lancaster

We. Communications

(415) 547-7054

hlancaster@wecommunications.com

Source: F5, Inc.