



NEWS RELEASE

# F5 Integrates F5 AI Guardrails with NVIDIA NeMo Guardrails to Extend Runtime Security for Enterprise AI Applications

2026-07-29

New integration helps deliver centralized policy enforcement across AI applications to drive faster enterprise AI deployments

SEATTLE--(BUSINESS WIRE)-- F5 (NASDAQ: FFIV), the global leader in delivering and securing every app and API, today introduced F5 AI Guardrails integrated with **NVIDIA NeMo Guardrails**, an AI-native solution that brings F5's enterprise-grade AI security capabilities to production AI applications.

Security remains a primary reason enterprises struggle to move AI from pilot to production. Organizations adopting multiple AI models and frameworks across hybrid multicloud infrastructures face fragmented security controls, inconsistent policy enforcement, and governance gaps that widen with every new AI application. Most AI frameworks include basic safety checks, but those controls are embedded inside individual applications rather than enforced centrally. Security and compliance teams are left stitching together visibility across tools and applications with no single place to inspect AI traffic, apply policy, or audit what happened.

F5 AI Guardrails integrated with NVIDIA NeMo Guardrails brings centralized, consistent security and governance across AI applications by separating security enforcement from the AI framework. Organizations can standardize on NVIDIA NeMo Guardrails while F5 AI Guardrails handles security inspection of prompts and LLM responses, applying enterprise policies centrally without requiring changes to AI applications. The result is a layered

architecture where the framework library, microservices, orchestration, and security each do what they do best and evolve independently.

“Enterprises do not have a shortage of AI applications. They have a shortage of consistent security and governance across them,” said Kunal Anand, Chief Product Officer at F5. “As AI moves into production, fragmented controls create risk, complexity, and delay. This integration gives security teams a unified point to inspect AI traffic, enforce policy, and govern AI applications across their environments, while giving developers the freedom to keep building.”

## Centralized security inspection for production AI

**F5 AI Guardrails** secures AI systems, inspecting prompts and responses between users and applications, and helping prevent prompt injection, PII exposure, data leakage, and malicious outputs. Security policies are applied centrally rather than built into each AI application, so organizations can enforce enterprise security requirements without modifying application code. Key capabilities include:

- Prompt and response inspection: Security policies are enforced on AI traffic in real time, helping detect and reduce prompt injection, unauthorized data disclosure, and harmful outputs before they reach users or downstream systems.
- Centralized policy enforcement: Security teams apply and manage AI security policies across applications from one place, without rewriting application code or modifying workflows.
- Consistent visibility across AI applications: A single view into AI application behavior across models, frameworks, and business units lets security teams monitor, audit, and act from one operational layer.
- Independent, layered architecture: F5 security inspection layers on top of the NVIDIA AI framework. Each can be updated, scaled, or replaced independently as requirements change.

This approach eliminates the tension between development speed and security requirements, enabling AI platform teams to move fast on orchestration while security teams maintain independent control over inspection, policy, and governance. With this solution, neither must wait for the other.

“NVIDIA NeMo Guardrails provides an open, programmable framework for applying safety and security policies to AI applications,” said Ash Bhalgat, Senior Director of AI Networking and Security Solutions, Ecosystem and Marketing at NVIDIA. “The integration with F5 AI Guardrails expands the range of protections customers can use to secure LLM prompts and responses as they move AI agents into production.”

## An open architecture that scales with AI adoption

F5 AI Guardrails integrated with NVIDIA NeMo Guardrails protects AI applications regardless of the underlying large

language model. Organizations can change models, add new AI applications, or expand across business units while maintaining consistent security inspection and policy enforcement. As AI adoption grows, security is designed to scale with it rather than constraining it.

F5 AI Guardrails integrated with NVIDIA NeMo Guardrails is generally available now.

## Supporting quotes

“As enterprises move AI into production, the infrastructure AI lives on becomes the competitive differentiator. Equinix provides the neutral, interconnected foundation where AI runs at scale, bringing together private infrastructure, multiple clouds, and a rich AI ecosystem. Our collaboration with F5 extends consistent security and policy enforcement across those distributed environments, enabling customers to scale AI with confidence while preserving performance, sovereignty, and choice,” said Maryam Zand, VP, Ecosystem Development, Equinix.

“While enterprises are rapidly adopting AI, scaling it remains a significant hurdle. The Red Hat AI Factory with NVIDIA is a co-engineered, enterprise-grade AI solution that helps organizations more seamlessly build, deploy, and scale AI workloads, wherever their business demands it. This latest integration of F5 AI Guardrails with NVIDIA NeMo Guardrails extends runtime protections to help customers deploy and scale AI with greater confidence using the Red Hat AI Factory with NVIDIA,” said Ryan King, Vice President, AI and Infrastructure Partners, Red Hat.

“Enterprise AI success depends on more than powerful infrastructure—it depends on establishing trust before applications reach production. Through WWT’s AI Proving Ground and our collaboration with NVIDIA, customers can validate AI architectures before deployment. The integration of F5 AI Guardrails with NVIDIA NeMo Guardrails further strengthens that approach by adding policy-driven AI governance and runtime protections that help organizations deploy agentic AI securely and with confidence,” said Chris Konrad, Global VP, AI and Cyber, WWT.

## Supporting resources

- **F5 technology alliance ecosystem**
- **F5 Labs CASI and ARS leaderboards**

## 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.

Dan Sorensen

F5

(650) 228-4842

**d.sorensen@f5.com**

Holly Lancaster

We. Communications

(415) 547-7054

**hlancaster@wecommunications.com**

Source: F5, Inc.