



Fastly Launches AI Firewall and AI Runtime Control to Secure and Scale AI

September 21, 2026

New capabilities govern AI requests, protect AI applications, and control agent interactions on Fastly's unified platform

SAN FRANCISCO--(BUSINESS WIRE)--Sep. 21, 2026-- Fastly, Inc. (NASDAQ: FSLY), a leader in global edge cloud platforms, today announced AI Runtime Control, AI Firewall, and new API Security capabilities designed to give organizations real-time visibility and control across their AI systems.

Expanding the Fastly for AI portfolio, these new capabilities help organizations govern AI model access and usage, protect the AI applications powering their business, and control how AI agents access enterprise APIs, all on the same Fastly platform already delivering the speed, security, and resiliency businesses depend on at global scale.

As enterprise AI shifts from experimentation to business-critical production, the focus turns to leveraging AI safely at scale. Organizations now need real-time control over model routing, spend, agent authorization, response optimization, and service reliability as traffic expands.

Fastly is seeing that shift directly. Across Fastly's network, machine-generated traffic crossed 50% in July and August this year, while AI traffic grew [6.5 times faster than human traffic](#) from January through May this year. Also, financial pressure is rising alongside this growth, with [McKinsey finding that 93% of organizations are exceeding their AI budgets](#).

Fastly extends its programmable, global edge platform to AI, giving organizations greater power and control to run at AI scale. Through the Fastly platform, consistent access, security, and routing policies can now be applied in the request path across vendors or model providers, so organizations can observe, route, secure, and govern AI activity as it happens, without giving up their choice of models or infrastructure.

"To effectively adopt coding agents and implement AI features without losing velocity or eroding existing business resilience, enterprises need control in production, at runtime, without delay, friction, or disruption," said Kelly Shortridge, Chief Product Officer at Fastly. "Fastly is extending the real-time control customers already trust for content delivery and software security to the models, applications, and agents shaping this new era of distributed systems. Organizations deserve the freedom to move fast with AI while staying safe, achieving innovation at scale without trading off resilience."

AI creates three critical paths that organizations can govern, protect, and control: applications calling models, users and systems interacting with AI applications, and agents calling enterprise APIs. At the same time, coding agents and AI-assisted development are creating new endpoints faster than they can be documented or reviewed. Today's launch extends Fastly's capabilities to address each of these:

- **AI Runtime Control (governing AI access and usage):** Routes model calls through a single endpoint across public and self-hosted providers, while virtual keys help protect underlying provider credentials. With real-time visibility into token spend, rate limiting, budget controls, and failover, organizations maintain model choice while enforcing policy centrally.
- **AI Firewall (protecting AI applications):** Mitigates LLM-based attacks, including prompt injection, delivering real-time protection at the edge. By evaluating prompts directly in the request path, organizations can identify and block malicious activity before it reaches targeted models.
- **API Security (controlling agents accessing enterprise APIs):** Enforces API contracts, applying consistent rules to agentic, agent-assisted, and conventional traffic. Organizations can observe or block non-conforming requests service by service, helping prevent agents from accessing unsupported operations, sending malformed requests, or acting outside established API boundaries.

Fastly for AI combines proven platform capabilities like Bot Management, ContentGuard, DDoS Protection, API Discovery, and Next-Gen WAF with new innovations. Today's launch extends the platform to control the impact of external AI traffic, govern internal AI activity, and protect autonomous AI agents acting on the organization's behalf.

Together, these capabilities help organizations build, secure, and deliver their applications, AI, and data on one platform, placing AI control on the same infrastructure already responsible for delivering and protecting their digital experiences.

Fastly's network runs at 622 Tbps of global edge network capacity (as of June 30, 2026), handling more than five trillion requests daily (as of March 31, 2026). Control over AI activity runs on that same powerful platform.

AI Runtime Control, AI Firewall, and API Security are now available. To learn more, visit fastly.com.

About Fastly, Inc.

Fastly's powerful and programmable edge cloud platform helps the world's top brands deliver online experiences that are fast, safe, and engaging through edge compute, delivery, security, and observability offerings that improve site performance, enhance security, and empower innovation at a global scale. Compared to other providers, Fastly's powerful, high-performance, and modern platform architecture empowers developers to deliver secure websites and apps with rapid time-to-market and demonstrated, industry-leading cost savings. Organizations around the world trust Fastly to help them upgrade the internet experience, including Reddit, Universal Music Group, and SeatGeek. Learn more about Fastly at <https://www.fastly.com>, and follow us @fastly.

Forward-looking statements

This press release contains "forward-looking" statements that are based on Fastly's beliefs and assumptions and on information currently available to Fastly. Forward-looking statements may involve known and unknown risks, uncertainties, and other factors that may cause its actual results, performance, or achievements to be materially different from those expressed or implied by the forward-looking statements. These statements include, but are not limited to, those regarding the ability of Fastly's products, including AI Runtime Control, AI Firewall, and its API Security capabilities to: give organizations real-time visibility and control across their AI systems; govern AI access and usage and provide real-time visibility; protect AI applications, including by identifying and blocking malicious activity and mitigating LLM-based attacks; and control how AI agents access enterprise APIs, including by enforcing API contracts; and the anticipated benefits to organizations from adopting Fastly's AI products. Except as required by law, Fastly assumes no obligation to update these forward-looking statements publicly, or to update the reasons actual results could differ materially from those anticipated in the forward-looking statements, even if new information becomes available in the future. Important factors that could cause our actual results to differ materially are detailed from time to time in the reports Fastly files with the Securities and Exchange Commission ("SEC"), including in our Annual Report on Form 10-K for the fiscal year ended December 31, 2025 and our Quarterly Reports on Form 10-Q. Copies of reports filed with the SEC are posted on Fastly's website and are available from Fastly without charge.

View source version on businesswire.com: <https://www.businesswire.com/news/home/20260921078405/en/>

Media Contact

Stacey Hurwitz
press@fastly.com

Investor Contact

Vernon Essi, Jr.
ir@fastly.com

Source: Fastly, Inc.