carahsoft

F5 AI Guardrails

Enable real-time protection for AI systems to safeguard data, enable compliance and maintain trust as you scale generative and agentic AI deployments.

Thank you for your interest in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies and a wide range of contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with F5, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.

To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com

Explore More Resources:
carah.io/f5-resources

Join Events & Webinars:
carah.io/f5-events

Discover Technology Solutions:
carah.io/f5-solutions

Learn About Procurement:
carah.io/f5-contracts

Connect With Our Team:
F5-Sales@carahsoft.com
(877) 95-F5GOV

SOLUTION OVERVIEW

F5 AI Guardrails

Enable real-time protection for AI systems to safeguard data, enable compliance, and maintain trust as you scale generative and agentic AI deployments.

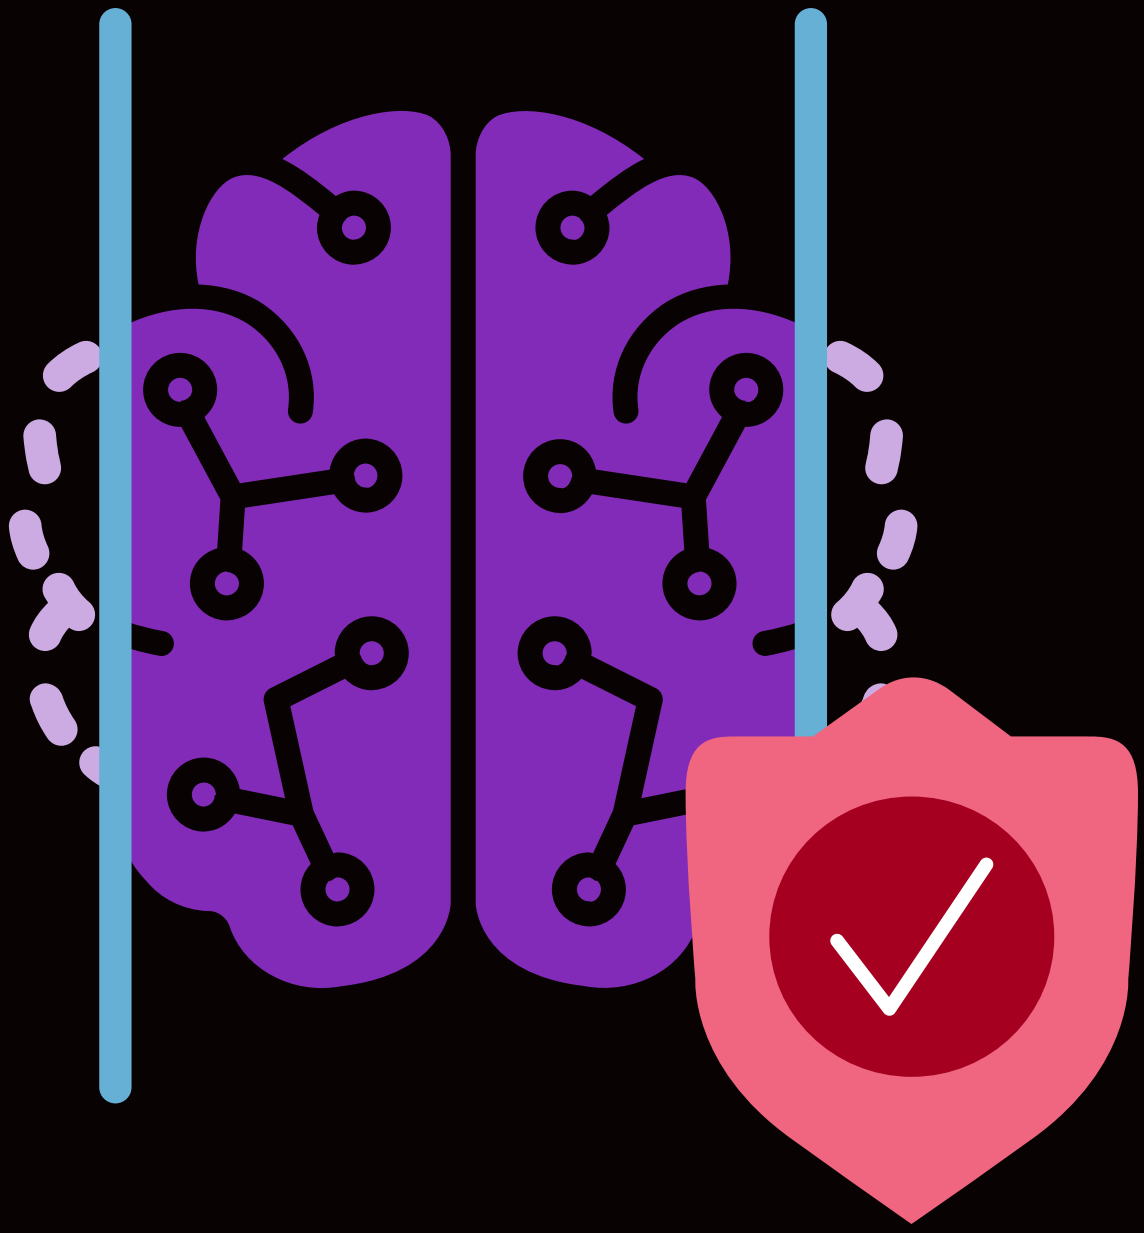

Key benefits

Protect AI systems from new runtime threats
Detect and block adversarial attacks such as prompt injection, jailbreaks, and data exfiltration during live AI interactions.

Enhance visibility and control
Monitor AI system behavior in real time with centralized dashboards and SIEM/SOAR integrations for major platforms.

Enable governance and compliance
Meet requirements for GDPR, HIPAA, EU AI Act, and more, with customizable and out-of-the-box policy enforcement and audit-ready logs.

Simplify observability
Achieve continuous visibility and traceability across all AI interactions, unifying performance, security, and compliance insights in one view.

Securing AI where it operates: at runtime

Enterprises are adopting AI at an unprecedented pace, embedding it into business-critical applications, decision systems and everyday workflows. This rapid adoption has introduced a new risk layer at runtime, where AI systems interact directly with users and sensitive data. Traditional application security tools cannot effectively secure and govern these interactions, leaving organizations exposed to novel threats such as prompt injection, data leakage, and unsafe model outputs.

Additionally, as AI scales across industries, regulatory expectations have grown equally complex with sprawling idiosyncrasies by region, jurisdiction, and industry. Frameworks like the EU AI Act, GDPR, and ISO/IEC 42001 require traceability, accountability, and real-time oversight of AI behavior. But most enterprises lack the visibility and control to keep pace with how quickly AI systems can create exposure once deployed. To maintain trust and compliance, organizations need a solution that enforces policies, protects data, and governs behavior in real time, directly where AI systems operate.

Enable real-time protection, compliance, and control with F5 AI Guardrails

F5® AI Guardrails secures AI applications, agents, and data where risk is greatest—at runtime. This solution continuously monitors AI interactions, enforcing enterprise policies to block prompt injection, data leakage, and harmful outputs before they reach end users. AI Guardrails operates independently of model providers, ensuring consistent protection across every public and private cloud, hybrid, and on-premises deployment.

The solution includes custom and out-of-the-box compliance guardrails that align with GDPR, HIPAA, and the EU AI Act, giving security and compliance teams confidence that AI operations meet regulatory requirements. Every enforcement action is explainable through outcome analysis, providing full transparency and audit-ready evidence of system behavior.

Strengthen your AI security posture with continuous improvement

When integrated with F5 AI Red Team, AI Guardrails becomes part of a continuous AI security lifecycle. AI Red Team insights identify vulnerabilities and inform policy updates in AI Guardrails, helping organizations rapidly adapt defenses as threats, models, and business needs evolve. Together, they help enterprises maintain compliance, trust, and resilience while enabling secure AI adoption at scale. Both products operate within the F5 Application Delivery and Security Platform (ADSP), unifying security for apps, APIs, and AI systems across hybrid and multicloud environments.

Key features

Detect and block adversarial threats
Apply real-time scanning and adaptive threat management controls at runtime to prevent prompt injection, jailbreaks, and other adversarial exploits.

Secure AI data
Deploy runtime data protection controls to detect and block the unauthorized transfer or misuse of sensitive information during AI interactions.

Adapt protection as needs evolve
Enable security teams to refine and update enforcement policies as applications, user behaviors, and data contexts change.

Deliver full transparency and auditability
Capture every enforcement event with outcome analysis and provide explainable, audit-ready visibility into AI decisions.

Unlock adaptive AI security for an ever-changing landscape

New threats require new security protections. F5 AI Guardrails equips enterprises to ensure AI security as models, users, and risks evolve.

Next steps

Explore how you can secure your AI with confidence
Learn more about F5 AI Guardrails.

Contact us
Reach out to discuss how F5 solutions can help you achieve your goals.

DC 10.2025 | JOB-CODE-1788985174