




F5 AI Red Team

Proactively test your AI systems—including models, applications, and agents—to find and fix vulnerabilities before attackers strike.



Key benefits

Expose vulnerabilities before deployment

Continuous, automated adversarial testing identifies risks long before production to reduce exploitation windows.

Validate resilience under pressure

Operational stress tests simulate latency overload, denial of service attacks, and resource drain to harden reliability and performance.

Deploy anywhere

Cloud, on-premises, or hybrid options support data sovereignty and compliance requirements.

Quantify and communicate risk

Every assessment is risk-scored against the Comprehensive AI Security Index (CASI) and Agentic Resistance Score (ARS) for objective evaluation, governance tracking, and executive reporting.

Seamlessly move from testing to protecting

Integration with F5 AI Guardrails enables the conversion of findings into runtime policies for live enforcement across environments.

AI systems are evolving faster than security teams can test them

AI security needs are changing rapidly—and many organizations are struggling to keep pace. Manual, weeks-long testing engagements can't match model and agent change rates. Frontier-model safety measures are opaque. And traditional AppSec tools miss runtime-layer threats like prompt injection, data exfiltration, jailbreak chaining, model distillation, and multi-agent privilege escalation.

To help enterprises keep up with new risks and secure their AI innovation, F5® AI Red Team exposes vulnerabilities before they become breaches through continuous, automated adversarial testing of apps, models, and agents against real-world attacks. Using an extensive AI attack database and red team agents, this solution delivers risk-scored reports that quantify AI security for insight-driven remediation. Integration with F5 AI Guardrails transforms testing results into continuous protection in an ever-evolving threat landscape.

Uncover vulnerabilities before attackers do

F5 AI Red Team combines three automated testing types—agentic resistance, signature attacks, and operational attacks—for full-spectrum validation. Agentic resistance tests run dynamic, multi-turn campaigns that emulate sophisticated real-world attackers and generate agentic fingerprints for transparent explainability. Signature attacks leverage tens of thousands of up-to-date prompts every month that keep testing aligned to emerging threat techniques, while operational attacks validate resilience under stresses such as crashes, resource exhaustion, or latency. Together, these methods deliver high-confidence vulnerability discovery across models, apps, and integrations.

Using this solution, security teams get prioritized remediation guidance in detailed reports that include successful malicious prompts, model responses, security scores, and severity classifications. Recurring campaign scheduling and CI/CD integration let organizations adopt continuous, automated testing, closing the gap between development and secure production rollouts. These insights also feed F5 AI Guardrails, enabling defenders to translate AI Red Team findings into runtime policies and protections rapidly.

Stay ahead of emerging AI risks

Automated, multi-step adversarial testing identifies vulnerabilities across AI systems using continuously updated, real-world attack prompts and agentic resistance tests that replicate the latest real-world adversarial behaviors. Every AI Red Team campaign produces measurable security scores to quantify exposure and resilience, ultimately guiding faster, data-driven remediation.

Additionally, agentic fingerprints provide natural-language explainability for every decision made during testing—showing exactly how vulnerabilities were exploited and why.

Key features

Agentic resistance

Multi-turn, intent-aware testing mimics human adversaries to reveal emergent vulnerabilities missed by static prompts.

Signature attacks

One of the largest and fastest-growing prompt databases—with 10,000+ new malicious prompts added monthly—helps keep campaigns aligned to the latest threats.

Automated scheduling and reporting

Recurring campaigns with severity scoring deliver continuous validation and executive-ready insights.

Agentic fingerprints

Step-by-step, natural-language explainability for each agentic decision creates an auditable trail that accelerates remediation and supports governance, risk, and compliance needs.

Operational attacks

Purpose-built testing—including latency overload, denial of service, and crash/resource exhaustion—helps ensure reliability and resilience.

Enterprise integration

Synergy with F5 AI Guardrails and the F5 Application Delivery and Security Platform unifies testing, defense, and observability across multicloud environments.

This transparency turns complex adversarial behavior into clear, auditable insight for faster remediation and defensible governance.

Validate AI resilience and performance

Organizations are looking to ensure AI reliability and stability under pressure. To enable this, AI Red Team delivers operational stress-testing that can identify performance bottlenecks, latency issues, and system weaknesses unique to AI workloads. This information helps security teams prevent downtime and ensure models perform securely at scale.

By taking advantage of AI Red Team, enterprises can:

- Simulate operational-level attacks such as latency overload, denial-of-service, and crash testing to assess system resilience under stress.
- Distinguish adversarial threats from operational risks, enabling a complete view of both exploitability and system resilience.
- Deliver comprehensive vulnerability reports with severity scoring to help prioritize remediation by business impact and operational criticality.
- Detect operational vulnerabilities that can cause cascading failures, degraded model accuracy, or instability during high-load conditions.

Accelerate AI security readiness

AI Red Team enables enterprises to shift from manual to continuous testing. Automated red-teaming delivers faster results, frees security experts from repetitive tasks, and ensures AI deployments remain secure while enabling compliance as threats evolve.

Plus, automated campaign scheduling provides always-on validation, ensuring models remain secure and resilient as configurations and threats evolve. Teams can generate audit-ready, explainable reports that support governance, risk, and compliance initiatives and provide executive-level visibility into testing coverage, vulnerability trends, and AI readiness posture.

Transform AI security from reactive defense to proactive assurance

F5 AI Red Team brings precision and scale to AI security testing by automating what once required manual expertise and weeks of effort. By continuously simulating real adversaries, organizations can move from reactive response to proactive assurance, reducing exposure windows and accelerating time-to-remediation.

The result: quantifiable, explainable insights for AI that's both innovative and secure.

F5 helps enterprises to securely scale, connect, and optimize AI workflows—maximizing performance and unlocking the full potential of AI.

