SPECTEROPS
Creators of BloodHound

DATASHEET

AI Red Team Enablement

Build AI-Native Offensive Security Capability

AI is changing how quickly adversarial capabilities are developed and adapted. SpecterOps researchers demonstrated that specialized offensive tooling that once took weeks to develop can now be created in hours with AI-assisted workflows. Security teams need to adapt at the same pace.

SpecterOps AI Red Team Enablement gives red teams the secure infrastructure, tooling, training, and practitioner expertise to incorporate AI into modern offensive security operations. Teams learn how AI changes adversary tradecraft while building the practical experience to test their own defenses against rapidly evolving techniques.

The result is capability that stays inside your organization, allowing your operators to continue testing, adapting, and improving as AI advances.

Amplify Your Team's Talent

SpecterOps helps teams apply AI-enhanced capabilities to their offensive workflows while maintaining the operational discipline required for enterprise security testing.

Pace with the Adversary

Develop and adapt offensive capabilities at the speed AI-assisted threats are evolving.

Cover More Ground

Accelerate research, tooling, analysis, and supporting tasks so operators can accomplish more within an engagement.

Prepare for Emerging Tradecraft

Exercise defenses against AI-assisted techniques before your security team encounters them during an incident.

Build Lasting Expertise

Develop AI-enabled offensive capability within your team that continues to grow between engagements.

AI-Enabled Tooling and Infrastructure

- SpecterOps works with red teams to apply AI-assisted development to real offensive operations, including custom tooling that can move from prompt to deployment in hours rather than weeks.
- Secure red team infrastructure provides isolated, per-project environments that deploy quickly and consistently, with role-based access controls and activity logging built in. Operators spend less time preparing infrastructure and more time conducting the engagement.
- SpecterOps Skills package practitioner expertise into reusable AI skills, plugins, agent definitions, references, and repeatable security workflows.

Tradecraft and AI Training

Hands-on training gives practitioners the foundation to build, evaluate, attack, and secure LLM and agentic systems. Courses cover areas including threat modeling, adversarial testing, reverse engineering, and practical defensive workflows.

Teams can extend that foundation through additional SpecterOps Tradecraft Academy courses in identity-driven offensive tradecraft, red team operations, Kubernetes, software supply chain security, and more. Training is available through in-person, virtual, private, and on-demand formats.

Applied AI Security Research

SpecterOps researchers work with frontier models to understand emerging offensive and defensive capabilities and translate those findings into practical security methods.

Our research spans AI-generated offensive tooling, LLM jailbreak testing, AI agent endpoint exposure, LLM-powered EDR analysis, and AI-assisted attack path discovery. These lessons inform our client engagements and the capabilities we make available to security teams.

Developed and refined through practical use, these resources give teams tested starting points for AI-assisted security work, including what to inspect, which dependencies and approvals matter, and how to validate results.

The SpecterOps Difference

Practitioner-Led Enablement

Learn from researchers and operators actively applying AI to offensive security, adversarial research, and real-world engagements.

Research Turned Into Tradecraft

Methods developed through SpecterOps AI security research become practical workflows, tools, skills, and training your team can use.

Built on Open Security Innovation

From BloodHound and Mythic to SpecterOps Skills and our latest AI research, we have a long history of turning offensive security expertise into capabilities the security community can build upon.

Capability That Stays With Your Team

AI Red Team Enablement builds internal expertise so your organization can continue testing and adapting as adversary capabilities evolve.

Getting Started

Simulate AI-enabled adversary techniques against enterprise defenses and the AI systems you've deployed.

See Your Environment the Way an Adversary Does

Simulate AI-enabled adversary techniques against enterprise defenses and the AI systems you've deployed.

Evaluate Frontier Models at Enterprise Scale

AI Cyber Ranges provide realistic environments for measuring advanced models across multi-step attack paths and interconnected technologies.

Build the AI-enabled offensive capability your security team needs to stay ahead of the adversary.

specterops.io/contact ›

©2026 Specter Ops, Inc.