SPECTEROPS
Creators of BloodHound

DATASHEET

AI Ranges for Enterprise Blue Teams

Flexible, Expert-Built Cyber Ranges for Testing AI With or Against Your Systems

Organizations want AI agents doing real operational work: administering infrastructure, handling sensitive data, and acting inside business-critical systems. Security leaders are being asked to approve that access without evidence of how the agent will use it.

SpecterOps AI Ranges for Enterprise Blue Teams give your team a contained enterprise environment for secure testing. Start from a prebuilt range or build one from a snapshot of your own environment. We bring the range design, adversary tradecraft, and guidance. Your team runs the tests, owns the analysis, and makes the calls.

Your range. Your tests. Our expertise.

A Range to Fit Your Specific Needs

AI behavior depends on the tools, identities, data, and trust relationships around it.
Choose the level of fidelity your evaluation needs and change it as your program matures.

Option	What You Get	Best For
Prebuilt Ranges	Ready-made enterprise environments modeled on representative AD, Entra ID, cloud, and hybrid architectures	Fast starts, model comparisons, team training
Custom Ranges	Environments tailored to your technologies, applications, and identity relationships, up to a full build from a snapshot of your production systems and data flows	Testing specific workflows and high-fidelity validation before production deployment

Every option is contained, repeatable, and reusable as models, permissions, and safeguards change.

How It Works

We provide the expertise and the environment. Your team drives the evaluation.

We Handle	You Control
Range design, build, and hosting	When and how you test
Prebuilt or custom environments from a snapshot of your systems	Criteria for successful testing
Expert-built scenarios grounded in real attack paths	Your results and analysis
Verified containment	Deployment decisions

Optional Add-Ons

Adversary stimulus. Need a realistic opponent? SpecterOps can supply the threat activity, through AI-driven or human red team operators, to pressure-test your agents and defenses.

Advisory support. Focused guidance around design, objectives, and interpretations that update as your models and permissions change.

The SpecterOps Difference

Flexible by Design

Prebuilt or fully custom. Scale fidelity up or down without changing platforms.

Your Platform to Use

Your team owns the testing and analysis. Run evaluations on your schedule, against your criteria.

Expert-Built Scenarios

Designed by practitioners who execute these attacks in real environments, so every scenario reflects a proven attack path.

Guidance When You Want It

Our experts advise on range design, evaluation objectives, and scenario selection.

Fully Contained

Containment is defined and verified before any AI runs, so testing stays inside the range.

Repeatable Baseline

The range becomes a reusable benchmark for future models, permissions, and integrations.

Built on Frontier AI and Adversary Experience

SpecterOps serves as the red team of record for multiple frontier AI companies and builds cyber ranges used by the UK AI Security Institute, among other AI and safety organizations. Our ranges all draw on that work and SpecterOps's years of executing attack paths in real enterprise environments.

Trusted by 300+

Fortune 500 and government agencies to inform the tradecraft behind every scenario.

From the Creators of

BLOODHOUND COMMUNITY EDITION

GHOSTWRITER

NEMESIS

MYTHIC

What Your Team Can Evaluate

Run legitimate workflows alongside adversarial scenarios, measure what matters to your deployment decision, and track how results change across matched trials as models, permissions, and safeguards evolve.

Enterprise Workflows

How agents retrieve sensitive data, execute code, administer infrastructure, and act across connected applications, and where they need approval or narrower permissions. Measure verified completion rate, time and cost per task, and human interventions.

Offensive Capability

Whether models can discover and traverse realistic attack paths, sustain multi-step reasoning, and adapt when controls interrupt them. Measure attack path success and failure rates across matched trials.

Abuse Resistance

Whether malicious instructions, compromised identities, exposed credentials, or manipulated tool outputs can redirect an agent toward unauthorized actions or data. Measure unauthorized actions attempted and carried out, sensitive test records accessed, policy violations blocked, approval gates respected, and legitimate tasks wrongly blocked.

Defensive Response

Whether an agent investigating an incident reaches the right conclusion, takes appropriate action, and avoids unnecessary access or disruption. Measure detection and containment time, disruptive actions, and restoration success.

Common AI Cyber Range Use Cases

Model Capability and Risk Evaluation

- Test whether autonomous agents can execute multi-step attacks, lateral movement, and privilege escalation
- Measure reasoning limits across long, complex attack chains
- Probe for unsafe outputs, policy failures, and unauthorized responses

Defensive Agent Development

- Give threat-hunting and response agents realistic activity, signatures, and attack scenarios to find
- Validate detection, containment, and recovery decisions

Pre-Production Agent Validation

- Confirm agents perform as intended before they receive real SOC or administrative access
- Check for leakage of sensitive data, PII, or restricted records
- Prevent excessive agency: unauthorized actions or drift beyond intended scope

Training and Experimentation

- Use the range as a safe playground for teams and agents to explore within agreed boundaries
- Record actions, tool calls, access attempts, and outcomes for review

---

Getting Started with AI Ranges for Enterprise Blue Teams

CISOs and AI governance leaders get an evidence base for deployment and risk decisions.
Engineering, IAM, security operations, and red teams get a repeatable environment to test, learn, and iterate.

Find out what your AI will do with access before you grant it.

Contact us at specterops.io/contact to scope your range.

Learn more at specterops.io/services ›

SPECTEROPS ©2026 Specter Ops, Inc.