

Uncover vulnerabilities in your enterprise AI apps

AT-A-GLANCE

Continuously uncover unknown security risks in your AI systems using Zscaler red teaming, powered by the most advanced attack database in the world.

Why Automated AI Red Teaming?










Manual AI red teaming is slow, inconsistent, and difficult to scale — creating major bottlenecks for security teams trying to keep pace with the rapid advancements and adoption of LLM-powered systems. Zscaler AI Red Teaming eliminates these challenges with the most advanced automated red teaming engine, which allows AI security teams and CISOs to evaluate AI systems faster, more frequently, and across a broader range of deployment scenarios.

With over 25 predefined probes and support for fully custom probes and dataset uploads, Zscaler AI Red Teaming enables organizations to tailor testing scenarios to their specific domains, regulatory environments, and use cases. Practitioners can simulate thousands of real-world attacks — across modalities like text, image, voice, and document uploads — within hours, surfacing gaps in security, safety, hallucination, and business alignment. By combining automation with customizability, Zscaler AI Red Teaming helps teams ship secure and compliant AI systems from day one — at enterprise scale.

Predefined Probes

 Text Variations  Image Variations  Voice Variations  Document Variations

Security:

Code Execution    
 Context Leakage    
 Data Exfiltration  
 Jailbreak  
 Manipulation  
 Phishing  
 RAG Poisoning 
 Web Injection 

Safety:

Bias    
 Cyber Threats    
 Fake News  
 Fraudulent Activities  
 Harmful Content 
 Illegal Activities  
 PII  
 Privacy Violation  
 Profanity  

Hallucination:

Paranoid Protection 
 Q&A 
 RAG Poisoning 
 URL Check 

Business Alignment:

Competitor Check  
 Intentional Misuse    
 Legally Binding  
 Off Topic    



Uncover Critical AI Vulnerabilities



Customize AI Risk Evaluations



Remediate Identified Risks



Export Shareable Testing Reports

Comprehensive AI Risk Assessments With Enterprise Scalability


GAIN VISIBILITY INTO YOUR AI'S RISK SURFACE

Evaluate your AI systems for weaknesses across all major risk categories — security, safety, trustworthiness & hallucination, and business alignment. With 25+ prebuilt probes and ongoing updates reflecting the latest attack techniques and strategies, Zscaler AI Red Teaming provides unmatched red teaming coverage for modern LLM-powered applications.




CUSTOMIZE PROBES & UPLOAD DATASETS

Go beyond standard testing by designing your own probes from scratch or uploading custom datasets with predefined attack prompts. Whether you're targeting specific behaviors or edge-case vulnerabilities, Zscaler AI Red Teaming gives you full control over your testing strategy — no matter how complex or domain-specific.



SIMULATE THOUSANDS OF ADVANCED ATTACKS


Run high-volume, domain-tailored adversarial simulations in minutes — at a scale that would be impossible to execute manually. Zscaler AI Red Teaming automatically adjusts to your system's rate limits and attack surface, surfacing critical security and safety risks faster and more reliably than manual red teaming assessments.





UNCOVER RISKS IN MULTIMODAL AI SYSTEMS

Zscaler AI Red Teaming is the first to simulate malicious inputs across text, image, voice, and document uploads — ensuring even your most advanced AI apps are secure. Review results through intuitive risk visualizations, and export full PDF reports to share with security and compliance teams, or other stakeholders.



Deploy Secure & Compliant AI Systems From Day One

Run advanced AI risk assessments pre- and post-deployment to uncover hidden weaknesses before they cause serious harm.

[BOOK A DEMO](#)

The screenshot shows a software interface for AI risk assessment. At the top, a header reads "Test Case 3 of 470". Below this, a table lists several test cases with columns for ID, Attack, Variation, Red Teamer, Strategy, and Result. The "Result" column for most cases shows "Passed", except for one which is "FAILED". To the right of the table, a "Comments" section contains a "Conversation 1" log. The log shows a "PROBE" message and a "CAR SALES" response. The "CAR SALES" response is a simulated AI interaction about finding a perfect car.

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange™ platform protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SSE-based Zero Trust Exchange™ is the world's largest in-line cloud security platform. Learn more at zscaler.com or follow us on Twitter @zscaler.

© 2026 Zscaler, Inc. All rights reserved. Zscaler™ and other trademarks listed at zscaler.com/legal/trademarks are either (i) registered trademarks or service marks or (ii) trademarks or service marks of Zscaler, Inc. in the United States and/or other countries. Any other trademarks are the properties of their respective owners.



**Zero Trust
Everywhere**