

AI Red Teaming

Emulate real-world attacks against AI systems



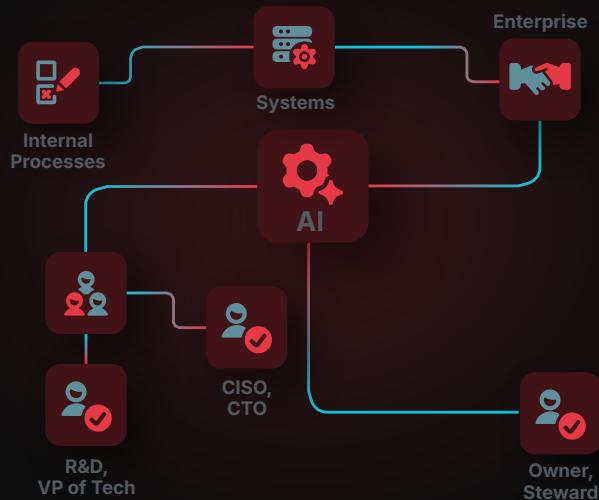
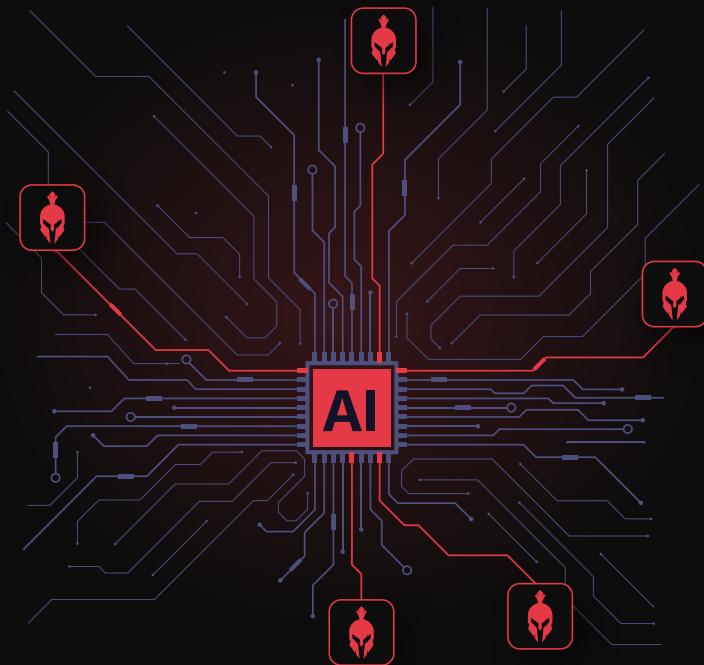
Solution Overview

Artificial intelligence introduces new classes of vulnerabilities that traditional security testing fails to uncover. Risks that can silently undermine enterprise security, brand integrity, and innovation investments.

Praetorian's AI Red Teaming service applies our industry-leading offensive security expertise to your organization's GenAI systems to simulate real-world attacks, assess defensive readiness, and deliver actionable recommendations for improvement.

Our engagements are designed to identify meaningful security gaps, avoid distraction from overhyped or low-impact issues, and demonstrate real-world exploitation scenarios that matter most to your business and R&D investment. This results in a technically rigorous assessment that delivers clear, prioritized insights.

Praetorian operators leverage advanced adversarial techniques such as RAG database poisoning, model theft, indirect prompt injection, and excessive agency to expose risk and measure true impact.



Who This Is For

- Enterprises embedding GenAI technology into products, systems, or internal processes
- Owners and stewards of organizational AI or GenAI risk
- CISOs, CTOs, Product Development VPs, and R&D Directors responsible for securing AI initiatives

Service Offerings

Praetorian offers three complementary engagement types, tailored to the maturity, complexity, and objectives of your GenAI program.

GenAI Penetration Test	GenAI Attack Path Mapping	GenAI Red Team Operation
<p>Guided by: OWASP Top 10 for LLM Applications</p> <p>Purpose: Designed for organizations integrating GenAI functionality at the application layer and beginning to formalize AI security testing.</p> <p>Consists of:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> A thorough review of the most salient LLM-application vulnerabilities <input checked="" type="checkbox"/> Demonstration of real-world exploitability and business impact <p>Most beneficial for:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Developers of LLM-enabled applications <input checked="" type="checkbox"/> Organizations already familiar with other OWASP Top 10 frameworks <p>Deliverables:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Executive summary <input checked="" type="checkbox"/> Vulnerability listing with CVSS 4.0 ratings <input checked="" type="checkbox"/> Proof of exploitation <input checked="" type="checkbox"/> Step-by-step reproduction guidance <input checked="" type="checkbox"/> Recommendations for improvement <p>Deliverables: 2-4 weeks</p>	<p>Guided by: MITRE ATLAS™</p> <p>Purpose: For organizations with more complex GenAI environments seeking to understand how vulnerabilities chain together to produce material risk and drive change.</p> <p>Consists of:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Exploitation of multi-stage vulnerability chains across GenAI systems <input checked="" type="checkbox"/> Targeting of specific objectives defined by the client to illustrate business impact <p>Most beneficial for:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Developers of complex GenAI systems <input checked="" type="checkbox"/> Organizations using GenAI beyond LLMs alone <input checked="" type="checkbox"/> Teams familiar with <i>MITRE ATT&CK®</i> <p>Deliverables:</p> <p>All items from the GenAI Penetration Test, plus:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Detailed attack-chain diagram <p>Deliverables: 4-8 weeks</p>	<p>Focus: Stealth and Evasion</p> <p>Purpose: A capstone engagement for organizations with significant AI investment—such as frontier-model developers or enterprises operating advanced GenAI infrastructure—seeking to evaluate detection and response under realistic adversarial conditions.</p> <p>Consists of:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Stealth & evasion red team operation targeting the organization's most sensitive AI assets <input checked="" type="checkbox"/> Targeting of the underlying AI infrastructure to steal sensitive models, proprietary information, and R&D investments <input checked="" type="checkbox"/> Optional follow-on purple team exercise for defensive tuning <p>Most beneficial for:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Frontier model developers <input checked="" type="checkbox"/> Orgs deploying complex GenAI-related infrastructure <input checked="" type="checkbox"/> Orgs operating extensive GPU compute capabilities <p>Deliverables:</p> <p>All items from the GenAI Attack Path Mapping, plus:</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Comprehensive attack narrative <input checked="" type="checkbox"/> Optional Purple Team report <p>Deliverables: 8-12 weeks</p>

MITRE ATLAS™ and MITRE ATT&CK® are a trademark and registered trademark of The MITRE Corporation.