ENTERPRISE EDITION · 2026

Striker *for LLM* Testing

Adversarial pressure testing that runs *before your attackers do.* Zero blind spots. Zero surprises.

Payload-driven simulation of real-world LLM attack patterns exposing exploitable weaknesses before a single release reaches production. Built for teams who ship AI at scale and can't afford a breach.

2,000+ real-world payloads · zero data egress

2,000+
ATTACK PAYLOADS
REAL-WORLD SCENARIOS

20+
ATTACKCATEGORIES COVERED

100%
CI/CD BLOCKING INTEGRATION

ATTACKCATEGORIES

Fifteen attack categories. One strike engine.

Strikermaps everypayloadto the OWASPLLMTop10 surfacingexploits your red team would find and the ones they'd miss.

PROMPTINJECTION
Direct & Indirect Injection
Single-turn and chained payloads embedded in user fields, tool outputs, or retrieved documents that hijack model directives.

ROLECONFUSION
Persona & Permission Escalation
Payloads that redefine model identity, elevate privileges, or suppress alignment behaviors through system-level impersonation.

MULTI-TURNCOERCION
Incremental Jailbreak Chains
Sequential conversation patterns that gradually erode safety guardrails across multiple exchanges invisible to single-turn scanners.

POLICYBYPASS
Guardrail Circumvention
Jailbreak-adjacent phrasing statistically correlated with policy violation across major model families and fine-tuned variants.

DATAEXFILTRATION
PII & System Prompt Leakage
Attack chains that coerce the model into revealing system prompt text, user PII, or confidential context via crafted outputs.

AGENTICEXPLOITATION
Tool & Workflow Abuse
Payloads targeting autonomous agent loops triggering unsafe tool calls, recursive execution paths, or unintended API side-effects.

STRIKEPIPELINE

Four stages. Fully automated.

EveryrunfiresinsideyourCIboundary blockingreleasesonexploitable findings before they ship.

01
Target Definition
Endpoint URLs, system prompts, model configs, and tool schemas normalized into a canonical strike profile.

02
Payload Dispatch
2,000+ categorised payloads fired sequentially and in parallel respecting rate limits and session state.

03
Response Scoring
Each model response graded against 15 exploit rubrics via a secondary judge model with audit evidence attached.

04
Report & Block
Findings ranked by severity. Critical exploits emit non-zero exit codes to halt CI/CD pipeline progression.

DEPLOYMENTOPTIONS

Your pipeline. Your thresholds.

No architecture changes. Striker drops into the stack you already run.

DEVELOPER · CI GATE

CLI & Pipeline Integration

Binarydrop-in forpre-commit hooks and CI gates. Block high-severity attack findings before staging.

- Pre-commit hook support
- GitHub Actions marketplace action
- Exit-code contracts for CI enforcement
- JSON + SARIF output for custom tooling
- Monorepo & multi-config support

PLATFORM · ENTERPRISE

Dashboard & API Server

Centralised findings with RBAC,trend analytics, and REST/gRPC API for full programmatic control.

- Role-based access control (RBAC)
- SSO via SAML 2.0 and OIDC
- Webhook alerts - Slack, PagerDuty
- Audit log trail for compliance
- Custom payload authoring interface

COMPLIANCEALIGNMENT **SOC 2 Type II** | **ISO 27001** | **NIST AI RMF** | **EU AI Act** | **OWASP LLM Top 10** | **GDPR**

> ***"Most security teams audit the model weights and ignore the instructions. Striker treats every system prompt, tool config, and conversation template as an attack surface — because that's exactly what an adversary does."***
>
> CoreLayer · SecurityResearchTeam · 2026

Request an enterprise evaluation.

Asolutionsengineer willscope Strikertoyourenvironment, wire it into your CI/CD pipeline, and deliver a full adversarial risk report within five business days.

REQUEST A DEMO

READ THE DOCS