AUTONOMOUS RED TEAM PLATFORM

GILDAL

Governed autonomous offensive security

Autonomous penetration testing that proves what's exploitable

A governed offensive-security platform: a domain-tuned AI red team **plans and executes the attack strategy itself**, keeps working even in fully self-hosted environments, and **generates a professional pentest report automatically** — while every risky action stays within an agreed scope and sealed inside an isolated environment.

- AUTONOMOUS KILL-CHAIN
- DOMAIN-TUNED AI MODEL
- SCOPE GUARDRAILS
- ZERO-ESCAPE ISOLATION
- COMPOUNDING EXPERIENCE
- AUTOMATED REPORTS

The guarantee

No false positives

Every vulnerability GILDAL reports comes with a working, reproducible proof of exploitation. **If it can't be exploited, it isn't reported** — so your team never chases ghosts.

Pentesting is slow and expensive. Automation loses control.

Organizations must test a wider attack surface more often, but skilled pentesters are scarce and costly. Plain scanners are fast, yet they bury teams in **context-free false positives**, and autonomous AI attackers are **hard to keep inside the lines**. GILDAL closes exactly that gap.

01 — COST

Expert-bound

Manual pentests are gated by headcount and scheduling — quarterly at best, never keeping pace with a shifting attack surface.

02 — NOISE

Context-free false positives

Legacy scanners match signatures; they never prove something is **actually exploitable**. Verification falls back on people.

03 — CONTROL

Ungovernable autonomy

Autonomous AI attacks risk drifting out of scope or leaking outward. Accountability and audit are hard.

04 — LOCK-IN

Cloud dependency

Relying only on frontier cloud LLMs brings cost, data-sovereignty, and isolated-environment constraints.

THE GILDAL ANSWER

A **governed AI red team** runs the kill-chain itself, but every risky action must clear a **scope gate** and stay inside a **sealed environment**. GILDAL reports only **confirmed vulnerabilities** — never false positives — each paired with a complete audit log.

From recon to proof, autonomously.

01 · RECON

Maps the attack surface

Explores your targets the way an attacker would — exposed services, endpoints and weak points.

02 · EXPLOIT

Chains real attacks

Autonomously chains vulnerabilities into working exploits — inside a sealed environment, strictly in scope.

03 · REPRODUCE

Confirms every finding

Each finding is independently re-tested and proven exploitable. No false positives.

04 · REPORT

Writes it up for you

A compliance-grade report with proof, impact and remediation — mapped to OWASP, CWE and CVE.

Every engagement sharpens the next

Findings and evidence accumulate automatically, so the platform reasons from everything it has seen before — getting faster and sharper over time, even in a fully **self-hosted deployment**.

Cloud

MAXIMUM CAPABILITY

Fully autonomous engagements with the highest detection capability — for formal, high-assurance assessments.

Self-hosted

FULL DATA SOVEREIGNTY

Runs entirely inside your environment — nothing leaves the perimeter. Built for regulated and air-gapped networks.

Domain-tuned proprietary AI model

Built for offensive security and grounded in MITRE ATT&CK, the CVE and CWE catalogs and a large corpus of real-world attack knowledge — so it recognizes attack patterns generic AI misses.

- MITRE ATT&CK
- CVE
- CWE
- EXPLOIT-DB

Autonomy and governance, together.

Scope guardrails

You define the allowed targets, tools, time-windows and blackout dates up front. Any action outside those limits is blocked the instant it's attempted — GILDAL **fails closed**, never open.

Zero-escape isolation

Every engagement runs inside a sealed environment. Traffic reaches the authorized targets and nothing else — **no path out, no pivot beyond scope**. Anything outside the boundary is refused.

Automated reporting

Explanations, remediation guidance, attack-surface analysis and an executive narrative are generated together into a **compliance-grade professional report** — automatically.

Threat-intel enrichment

Real NVD CVE · MITRE ATT&CK · CWE mapping with automatic **OWASP/CWE classification**, so every finding aligns to standard frameworks for executives and engineers alike.

Multi-target campaigns

Assess many targets sequentially or in parallel and generate a campaign-level report. One target failing **preserves the rest**; failures are reported separately.

Console & client portal

A real-time **operations console** for your security team — sessions, findings, attack paths, live activity — plus a **client portal** for delivery, and an API when you need to automate.

Capability, made concrete in the deliverable.

Red Team Assessment Report

GILDAL AUTO-GENERATED
OWASP · CWE · CVE MAPPED

2	3	5	4
CRITICAL	HIGH	MEDIUM	LOW

F-001	Remote Code Execution — Tomcat manager (uid=0 root)	CONFIRMED
F-002	Lateral Movement — cross-host SSH pivot (deploy)	CONFIRMED
F-003	SQL Injection — authenticated endpoint	VERIFIED

Multi-format output

HTML · PDF · DOCX from one button.

KISA 28-item checklist

Auto-scored to the Korean assessment standard.

Complete audit log

Every action recorded — a full audit trail.

Time to light up the command center.

GILDAL — GOVERNED AUTONOMOUS OFFENSIVE SECURITY

REQUEST A DEMO →