



OFFENSIVE SECURITY PLATFORM — RESTRICTED, VETTED ACCESS ONLY

Find what an attacker would find — before they do.

Reconnaissance, exploitation, and validation run against your domains, web apps, and IP ranges, and turned into evidence your team can act on the same day — not a CVE list to triage for a week.

45+

independent scan modules feeding
one report

ON-PREM

deployable on your own
infrastructure, no mandatory vendor
API

API + MCP

every action scriptable — and
callable directly by an AI agent

EXHIBIT 01

Security tooling wasn't built for how fast attack surface changes.

A new subdomain, a reopened port, a redeployed endpoint — your attack surface moves every day. Most security programs still check it once a quarter, with three different tools that don't talk to each other.

01

Fragmented tools

A port scanner here, a DAST tool there, a separate threat-intel feed, a spreadsheet to track what a pentester found eight months ago. Nothing shares context.

02

Annual, not continuous

A once-a-year pentest engagement is a snapshot. It says nothing about the subdomain that went live last Tuesday, or the port reopened after the last patch cycle.

03

Scores, not proof

A severity score with no reproducible evidence makes every finding a research project before a developer can even confirm it's real, let alone fix it.

Vulnix0 replaces that stack with one engine that re-runs on schedule, correlates findings across every module, and ships proof — not a guess — with every result.

02

EXHIBIT 02

Hunt. Exploit. Prove. Repeat.

Four stages, not a single scan. Each one feeds the next, and the loop restarts the moment your infrastructure changes — not on a yearly calendar entry.

01**Discover**

Reconnaissance agents enumerate domains, subdomains, APIs, cloud assets, and exposed credentials, then build a map of how they connect — wordlists, passive CT logs, and certificate SAN mining together, not just one source.

02**Exploit, under policy**

Each candidate finding is actually attempted — not inferred from a banner — within the exact module selection and aggressiveness ceiling you set per scan.

03**Prove & prioritize**

Every finding ships with reproducible request/response evidence, CVSS, CWE, and MITRE ATT&CK mapping. A SQLi on a staging form and one on your auth endpoint get very different priority here.

04**Repeat, on change**

A new subdomain, an opened port, or a deployed endpoint triggers re-testing on its own. Continuous diffing surfaces exactly what changed since the last scan.

REAL PENETRATION TESTING, NOT JUST A SCANNER

	TRADITIONAL PENTEST ENGAGEMENT	VULNIX0, CONTINUOUSLY
Cadence	Once or twice a year, scheduled weeks out	Every time your infrastructure changes
Technique depth	Manual, tester-dependent, varies by engagement	The same 45+ technique classes, every run
Cost per engagement	Day-rate consultants, scoped and invoiced per test	Fixed platform cost, unlimited re-runs
What you get	A PDF, weeks after the window closes	Live findings with reproducible proof, same day



EXHIBIT 03

One scan. Every angle.

Every module below runs independently and feeds the same evidence engine — pick exactly what to run per scan, from a five-minute recon pass to a full authenticated pentest.

ATTACK SURFACE & RECON

- Port scanning with service fingerprinting
- Subdomains: wordlist + passive CT logs + certificate SAN mining
- Subdomain takeover & lookalike-domain monitoring
- WHOIS, DNS records, DNSSEC validation, zone hygiene
- Cloud storage exposure: S3, GCS, Azure Blob, DO Spaces

NETWORK & INFRASTRUCTURE

- Native TLS cipher/chain audit & cert-expiry alerts
- Anonymous LDAP, weak RDP, exposed Docker/etcd/K8s
- Exposed MongoDB, Elasticsearch, Memcached, Telnet
- WAF/CDN fingerprinting, HTTP method/WebDAV audit

WEB & API SECURITY

- Crawl-driven DAST: injection, SSRF, IDOR/BOLA, open redirect
- API testing from an OpenAPI/Swagger spec
- GraphQL introspection & playground exposure
- Request smuggling, Host header injection, credentialed CORS
- JWT/session weaknesses & secrets-in-response scanning

RESILIENCE TESTING

- Configurable Stress Test under real load
- Bounded DoS-resilience & open-resolver checks

IDENTITY & EMAIL TRUST

- Default-credential & login-lockout testing
- Account enumeration on signup/reset flows
- SPF/DMARC/DKIM spoofing-risk verdict, DKIM key strength
- MTA-STS & TLS-RPT mail-transport hardening

PROGRAM & COMPLIANCE

- Continuous attack-surface diffing between scans
- OWASP Top 10 / PCI-DSS scorecard from real findings
- Scheduled Automation, projects, full audit log



EXHIBIT 04

Not “a scan.” The same techniques an attacker runs.

Every module maps to real, named technique classes — CWE-tagged, OWASP-mapped, MITRE ATT&CK-mapped — grouped below the way an actual intrusion unfolds: find the surface, break in, then see how far it goes and whether it holds under pressure.

01 — RECON & DISCOVERY	02 — EXPLOITATION	03 — VALIDATE & STRESS
Map everything before touching anything.	Actually attempt it — not inferred from a banner.	Confirm impact, then test if it holds up.
Subdomain Takeover	SQL Injection	Login Lockout Bypass
Typosquat Domains	Cross-Site Scripting	Account Enumeration
Cloud Bucket Exposure	SSTI	Secrets-in-Response
Open DNS Resolver	Command Injection	Anonymous LDAP / Weak RDP
DNSSEC Gaps	SSRF	Exposed DB (Mongo/Elastic)
Cert/SAN Mining	IDOR / BOLA	Stress Test
Shadow Sitemap Disclosure	Open Redirect	DoS-Resilience Check
	Request Smuggling	Attack-Surface Diffing
	Host Header Injection	
	Default Credentials	
	JWT Alg/Secret Weakness	
	Credentialed CORS Bypass	

Every finding carries a CVSS score, CWE ID, OWASP Top-10 category, and MITRE ATT&CK technique — the same taxonomy a SOC or a pentest report already uses, so nothing needs re-translating before it reaches a ticket.

05

EXHIBIT 05

AI doesn't just read the report. It can run the scan.

Two different AI surfaces, two different jobs: one lets an agent drive the platform end-to-end; the other turns a dense findings list into something a non-specialist actually reads.



Agent-driven scanning (MCP)

A native MCP server exposes the whole engine as callable tools: start a scan with the exact same module/config options as the UI, poll for results, and pull a diff against the asset's previous scan — all from a prompt, no API wrapper to write.

“Ask OneFirewall AI” (grounded summarization)

On every completed scan, one click generates a plain-English executive summary of what was actually found — and a one-paragraph explanation for any single finding — built strictly from that scan's own data, not a generic canned write-up.

“Latest scan added 3 findings, resolved 1, opened 2 ports, and changed the TLS certificate — here's what actually needs attention first.”

Illustrative output shape from the asset-change comparison an agent can request after any new scan.

06

EXHIBIT 06

Built for how security actually gets fixed.

Not another dashboard full of scores nobody reads — a tool that finds real issues, proves them, and keeps checking.

Runs where your data already lives

Deployable entirely on-prem: local GeoIP/ASN databases, no mandatory third-party API for core scanning, nothing leaves your network unless you choose an optional integration.

Evidence, not guesses

Every finding ships with reproducible request/response proof, CVSS, CWE, and MITRE ATT&CK mapping — verify it in minutes, don't re-run the scan.

Everything a human can click, an agent can call

A full REST API and a native MCP server expose the same engine to scripts, CI pipelines, and AI agents — configure every module exactly as in the UI.

One scan becomes a program

Projects, scheduled Automation, an audit log, continuous diffing, and a compliance scorecard turn a point-in-time scan into an ongoing practice.

Red-team rigor, every day

The same exploitation techniques a red team runs in a two-week engagement — attempted, not inferred — run continuously and automatically: authorized, bounded, and logged every time, not scoped to a once-a-year window.

Access is vetted, not self-serve — request a review for your environment.

[vulnix0.com](#) · [API Docs](#) · [MCP Server](#)

