CONTINUOUS THREAT EXPOSURE MANAGEMENT · VALIDATION

Pythia — continuous threat validation

The oracle for your CTEM cycle.

Scanners and feeds tell you what could be exposed — against the world's stack, not yours. Pythia validates every emerging advisory, KEV entry and PSIRT bulletin against **your actual environment**, every day, and turns the result into the **two or three actions that matter this week**. Privacy-defensible by construction.

01 The gap in every CTEM program

Continuous Threat Exposure Management runs in five phases. Most vendors sell you one of them. The Vulnium program covers the cycle end to end — specialists and Pythia divide the work — and the phase most programs still run by hand, **Validation**, is automated and evidenced daily.

PHASE 01

Scoping

Your environment profiled with our consultants — vendors, assets, controls.

SERVICES + PYTHIA

PHASE 02

Discovery

Assessments set the internal baseline; Pythia watches the external signal daily.

SERVICES + PYTHIA

PHASE 03

Prioritization

Every threat rated against your stack, not the world's — with reasoning.

PYTHIA

PHASE 04

Validation

Pythia validates continuously; our red team proves exploitability.

PYTHIA LIVES HERE

SERVICES + PYTHIA

PHASE 05

Mobilization

A tracked P1/P2/P3 checklist; our engineers help close it.

SERVICES + PYTHIA

02 Why validation is the bottleneck

95%
of threat news is irrelevant to any given environment — but someone still has to prove it

8–10 h
of analyst time per week spent hand-correlating "does this apply to us?"

~60 s
per Pythia intelligence run: collect, enrich, validate, recommend, distribute

<$1/day
total operating cost — model calls metered and reported per run

03 Validation has two layers — the program runs both

Proving a threat is real takes both continuous analysis and human adversaries. Point tools offer one; the Vulnium program runs both.

Adversarial validation VULNIUM SERVICES

Our pentesters and red team prove exploitability by hand: point-in-time engagements that confirm an attacker can actually reach and abuse an exposure — and that your detection fires. OSCP/OSWE/OSWP-certified, OWASP/OSSTMM/PTES/NIST SP 800-115 methodologies.

```
// deep · human · definitive — a snapshot in time
```

Threat validation PYTHIA

Between engagements, Pythia validates the external threat axis daily: is this new CVE exploited in the wild, does it hit a vendor you actually run, do your stated controls already blunt it? Each verdict carries reasoning and the exact assets affected.

```
// continuous · company-aware · always current
```

04 How a run works

Every run follows the same auditable pipeline. Stages that touch your company profile execute on a **local model only** — it never leaves your machine.

STAGE	RUNS ON	WHAT IT DOES
Collect	SCHEDULED	CISA KEV & Alerts, vendor PSIRT feeds, OTX, abuse.ch, MITRE ATT&CK, research blogs and social aggregators.
Enrich	PUBLIC LLM	Summarize each item; extract CVEs, ATT&CK TTPs, threat actors, affected products. Article content only — no profile data.
Validate	LOCAL LLM	Rate every item High / Medium / Low / N-A against your environment profile, with written reasoning and the exact assets affected.
Recommend	LOCAL LLM	Consolidate validated High/Medium items into a prioritized P1/P2/P3 action checklist with concrete steps and source links.
Distribute	YOUR CHANNELS	Microsoft Teams card, leadership dashboard, printable board briefing pack. Free-form Q&A with cited answers on demand.

05 Enforced in code, verifiable in audit

Fail-closed privacy boundary

Your environment profile — vendors, assets, controls — can never reach a hosted model by accident. A canary guard rejects any profile-bearing prompt bound for a public provider, and routing rules that configuration cannot override pin profile stages to the local model.

```
raise PrivacyViolation # profile × hosted provider
```

Tamper-evident audit trail

Every model call is recorded: provider, model, tokens, cost, and profile-bytes-sent. The audit page proves the invariant with a single number your legal team can verify themselves — evidence ready for SOC 2 and cyber-insurance review.

```
✓ Verified — 0 bytes of profile sent to any hosted provider
```

06 What your team gets

Executive home

Today's validated High/Medium items, KPI cards and run status — ready when leadership opens the laptop.

Grounded verdicts

Every rating tied to a public source and a specific profile match. No "where did this come from?"

Plain-English Q&A

"Does CVE-X affect us?" answered in seconds, with citations — on a local model.

Action checklist

P1/P2/P3 with concrete steps, live status tracking, and a printable board briefing pack.

Injection hardening

Authoritative domains are evidence; everything else is treated as adversarial input.

Remediation plans (ROADMAP)

Env-specific mitigate-vs-patch decisions, HA-safe sequences and compromise checks.

The Vulnium CTEM Program — how engagements start

One owner for the whole cycle: Pythia validates continuously, Vulnium specialists prove and close — evidence at every phase.

STEP 1 Profile & deploy

We model your environment (vendors, assets, controls) and deploy Pythia on your infrastructure — data stays on your box.

STEP 2 Continuous validation

Daily runs validate the landscape against your profile; leadership gets the checklist, analysts get their week back.

STEP 3 Adversarial proof

Scheduled Vulnium pentest / red-team engagements prove exploitability on what Pythia surfaces — loop closed.