

GHOST DEFENSE

Grounded Handling Of Sourced Threat-intel

VULNERABILITY EXPLOIT PRIORITY RANKER

V E P R

How-To Guide

Tool 06 · Build 2026.10.16-B · 16 October 2026

Rank vulnerabilities by what attackers can actually weaponize —
on the assets where it matters.

For defensive and authorized security use only · DEFEND ALWAYS.

Not a CISA product. Carries no CISA endorsement.

1. WHAT VEPR DOES

You have more vulnerabilities than you can patch this week. VEPR answers one question: which ones can an attacker actually use against you right now?

It takes a list of CVEs — pasted, or imported straight from a scanner — and cross-references four independent public sources to produce three separate verdicts on each one.

Verdict	Question it answers	Who decides
Priority (P1–P4)	Is there a working, public exploit for this, and is it confirmed exploited in the wild?	The tool, from public data
SSVC	CISA's decision tree: act, attend, track*, or track	Tool + your stakeholder context
BOD 26-04	The remediation deadline the directive assigns	Tool + your exposure and impact answers

They are deliberately kept separate and all three are shown. They can disagree, and when they do, the disagreement is information — it usually means a factor only you can supply is missing.

What VEPR does not do

- **It does not discover assets.** It does not scan. It has no network access to your estate and never touches your assets. You bring the findings; it ranks them.
- **It does not upload your data.** It is a browser page. Nothing you paste or import leaves your machine — see section 10.
- **It does not decide for you.** A patch is a decision about your environment. VEPR narrows the field and shows its work; the call is yours.

THE CORE RULE VEPR distinguishes "no data" from "no finding" everywhere. A blank is never shown as a clean result. If a check has not run, it says NOT CHECKED. If it ran and found nothing, it says NONE FOUND with the date. Those two mean completely different things and the tool will never let them look alike.

2. QUICK START (TWO MINUTES)

1. **Open** Open the tool. No install, no account, no network setup.
2. **Load your CVEs.** Paste your CVE list into the big box — one per line, or comma/space separated. Or click IMPORT SCAN to load a Nessus .nessus, Qualys CSV, or nmap XML export directly.
3. **Click LOAD EXPLOIT INTEL.** This is the step people skip, and skipping it caps every result at P3. It pulls Metasploit, Exploit-DB and Nuclei so the tool knows which CVEs have working public exploits.
4. **Click RANK NOW.** The ranking appears, worst first.

That gives you a usable P1–P4 triage list. The remaining sections cover the two richer verdicts — SSVC and BOD 26-04 — which need a small amount of input that only you can provide.

3. STEP ONE — GETTING YOUR VULNERABILITIES IN

Method	How	When to use it
Paste CVEs	Type or paste into the main box	A short list, or a handful you are investigating
Import scan	IMPORT SCAN → Nessus .nessus, Qualys CSV, or nmap XML	Best option. Findings carry host, port and service, so the ranking knows which assets are affected
Load live KEV	LOAD LIVE KEV pulls the full CISA catalog	Triaging the whole KEV catalog rather than your own findings
Upload KEV file	UPLOAD KEV FILE takes the KEV JSON or CSV	Air-gapped or restricted networks where the fetch is blocked

IMPORT THE SCAN IF YOU CAN A pasted CVE list tells you a vulnerability is dangerous. An imported scan tells you it is dangerous AND sitting on six specific hosts. Everything downstream — the P1 banner, the report, the asset counts — gets sharper when asset context is present.

The Try chips

The chips under the input box (F5 BIG-IP Set, Critical 2024, Ransomware CVEs, Network Devices, Load Full KEV) load worked examples. The active chip turns white — and it stays white only while the text still matches exactly. Edit one character and the highlight drops, so the chip can never claim you are looking at a set you have since changed.

4. STEP TWO — LOAD EXPLOIT INTELLIGENCE

This is the single highest-value click in the tool, and the one most often missed.

Until you load it, the tool does not know whether a working exploit exists. It refuses to guess: every KEV entry is capped at P3 and the EXPLOIT AVAILABILITY column reads NOT CHECKED. That is the tool telling you a check has not run — not telling you that you are clean.

Feed	What it proves	Source
Metasploit	A weaponized, point-and-click module exists	rapid7/metasploit-framework
Exploit-DB	Working exploit code is published	gitlab.com/exploit-database/exploithub
Nuclei	A detection template exists (detection, not exploitation)	projectdiscovery/nuclei-templates

Freshness

Intel is cached locally and re-used on your next visit, but it carries a 24-hour service level. Past that, the tool marks it stale and withholds SSVC outcomes rather than scoring you against last week's picture. Re-click LOAD EXPLOIT INTEL to refresh.

IF EXPLOIT-DB NEVER APPEARS Metasploit and Nuclei are served from GitHub, which browsers can fetch directly. Exploit-DB is served from GitLab, which does not send the headers a browser needs, so that one feed alone has to travel through a proxy. If you consistently see METASPLOIT and NUCLEI badges but never EXPLOIT-DB, the proxy hop is failing — the other two feeds are still valid.

5. STEP THREE — READING THE PRIORITY RANKING

	Meaning	Act within
P1 — PATCH IMMEDIATELY	In CISA KEV AND a weaponized exploit exists in Metasploit, Exploit-DB or Nuclei. Anyone can fire this today.	24–48 hours
P2 — PATCH THIS WEEK	In CISA KEV AND a public proof-of-concept exists. Moderate barrier to weaponization.	7 days
P3 — PATCH THIS MONTH	In CISA KEV, but no public exploit found. Confirmed exploited in the wild by someone.	30 days
P4 — ROUTINE	Not in CISA KEV. No confirmed in-the-wild exploitation.	Normal cycle

P2 IS CURRENTLY UNREACHABLE — KNOWN GAP P2 depends on a public proof-of-concept feed that is declared in the tool but not yet populated by any source. Until it is wired, no finding will be assigned P2; KEV entries with a PoC but no weaponized module will appear as P3. This is disclosed rather than hidden. Treat P3 as "KEV, exploit status not fully characterised".

The EXPLOIT AVAILABILITY column

Badge	What it means
METASPLOIT / EXPLOIT-DB / NUCLEI	Named database carrying an exploit for this CVE
NONE FOUND (date)	All feeds were searched on that date and none carried it. Private and nation-state tooling is not observable — this is not proof none exists.
? NOT CHECKED	Exploit intel has not been loaded. No conclusion has been drawn.

6. STEP FOUR — SSVC AND YOUR STAKEHOLDER CONTEXT

VEPR implements CISA's full four-factor SSVC decision tree, transcribed cell-for-cell from Tables 8 and 9 of the CISA SSVC Guide (November 2022). All 36 rows are unit-tested against the published values.

Three of the four factors the tool can derive. The fourth — Mission & Well-being — is a property of your organisation, and no feed can supply it.

Why the box is white

Throughout VEPR, white marks input that belongs to you. Green is what the tool derived. The STAKEHOLDER CONTEXT block is white because those two dropdowns are yours alone:

- **Mission prevalence:** minimal, support, or essential — how much your mission depends on the affected systems.
- **Public well-being impact:** minimal, material, or irreversible — what happens to people if it fails.

What happens if you leave it unset

The tool does not pick a default, because any default would invent a fact about your environment. Instead it evaluates the row at all three levels and shows a RANGE — for example TRACK-ACT ? — with the outcome at each level listed in the detail panel. Rows are sorted by their worst case, so nothing dangerous gets buried by uncertainty.

If the range collapses to a single outcome, the answer is certain without any input from you, and the tool says so.

7. STEP FIVE — BOD 26-04 REMEDIATION TIMELINES

CISA Binding Operational Directive 26-04, "Prioritizing Security Updates Based on Risk", assigns a remediation deadline from four yes/no variables. VEPR implements its decision table and shows the resulting tier beside the SSVC verdict.

Tier	Requirement
3 DAYS + TRIAGE	Remediate within 3 days AND perform forensic triage to determine whether the asset is already compromised. Patching does not evict an attacker who is already in.
3 DAYS	Remediate within 3 days.
14 DAYS	Remediate within 14 days.
60 DAYS	Remediate within 60 days.
FIX ON UPGRADE	Remediate at the next scheduled major upgrade or rebuild of the asset.

The four variables and where each answer comes from

Variable	Source, in precedence order
In KEV	The live CISA KEV catalog. Fully automatic.
Publicly Exposed	YOU ONLY. Exposure is a property of the asset, not the CVE, so no feed can supply it. Set an estate default in the BOD bar, or answer per-CVE in the detail panel.
Automatable	Your override → CISA Vulnrichment → inferred from public tooling → unknown
Technical Impact	System owner override → CISA Vulnrichment → CVSS vector → CVSS score → unknown

Click LOAD CISA VULNRICHMENT to fetch CISA's own published answers for Automatable and Technical Impact. Roughly half of all CVE records are unenriched; when one is, the tool says so rather than quietly falling back.

The system-owner override on Technical Impact

Only the system owner knows the real impact. A CVSS vector is a vendor's abstract claim about a component in isolation; CISA's enrichment is necessarily generic across every deployment. Whether a flaw yields total control of your component can turn on local configuration — whether the service runs privileged, whether the affected module is even enabled.

So the Technical Impact row carries Total / Partial buttons, and your answer outranks both CISA and the CVSS vector. The panel records which source supplied the value, so a reviewer can always see whether a tier rests on a published fact or on an operator assertion.

TECHNICAL IMPACT IS NOT MISSION CRITICALITY TOTAL means the flaw yields full control of the vulnerable component. PARTIAL means anything less. It is not a measure of how important the system is. BOD 26-04 has four variables and none of them is mission — deliberately, so the same flaw carries the same deadline at every agency. Marking a box TOTAL because it matters, rather than because the flaw grants full control, inflates the tier and the output stops being the directive. Mission weight has its own home: the SSVC Mission & Well-being control, which feeds SSVC and deliberately does not feed this timeline.

When no tier is shown

If Automatable or Technical Impact is unknown, VEPR assigns no tier at all and tells you which variable is missing. If only exposure is unset, it shows the answer as a span — "14 DAYS if not publicly exposed, 3 DAYS if it is" — displayed at the worse case until you set it.

8. STEP SIX — PRODUCING THE REPORT

The print/PDF report leads with P1 and P2 findings and carries the full methodology and provenance block: which factors were derived, which were analyst-set, which were operator-supplied, and the edition of the CISA guide applied.

- Ranges print as ranges. A printed verdict outlives the screen it was read on, so an unresolved SSVC outcome is never flattened into a single confident answer on paper.
- Headline count tiles state what was not evaluated — "N of M unresolved because Mission & Well-being is not set" — instead of implying a clean zero.
- Analyst-set values are marked "(analyst)" so a reviewer can separate your judgement from published data.

9. WHERE THE DATA COMES FROM

Source	Supplies	Refresh
CISA KEV catalog	Confirmed in-the-wild exploitation, due dates, ransomware flag	Live, 24h SLA
Metasploit Framework	Weaponized module availability	Live, 24h SLA
Exploit-DB	Published exploit code	Live, 24h SLA
Nuclei templates	Detection template availability	Live, 24h SLA
CISA Vulnrichment	CISA-published Automatable and Technical Impact	On demand
NVD	CVSS scores and vectors, vendor/product backfill	On demand
FIRST EPSS	Exploit prediction score	On demand

10. DATA HANDLING

- **Nothing you load leaves your machine.** Scan files are parsed in your browser. No CVE list, hostname, IP or scan file is uploaded anywhere.
- **Outbound requests go out** Only to fetch public catalogs — KEV, Metasploit, Exploit-DB, Nuclei, NVD. These are one-way reads of public data; your findings are never included in a request.
- **Settings stay local.** Your stakeholder context, exposure settings and overrides persist in browser local storage under ghost06_* keys, on your machine only. Clearing site data clears them.
- **It works offline.** The tool runs entirely from a single HTML file. It works from a USB stick on an isolated network — load the KEV JSON via UPLOAD KEV FILE and the ranking still works, without the exploit-availability layer.

11. TROUBLESHOOTING

Symptom	Cause and fix
Everything shows NOT CHECKED	Exploit intel is not loaded. Click LOAD EXPLOIT INTEL.
Every CVE comes back "NOT IN KEV" at P4	The KEV catalog is empty. Click LOAD LIVE KEV, or upload the KEV file. The tool withholds ranking on an empty catalog rather than reporting a false all-clear — if you see this, the fetch failed.
No exploit feed will load, page opened by double-clicking the file	A page opened from file:// sends an origin that proxies reject. Serve the folder over HTTP instead: from that directory run <code>python3 -m http.server 8000</code> and open <code>http://localhost:8000/</code> — or just use the deployed site.
METASPLOIT and NUCLEI appear, EXPLOIT-DB never does	Exploit-DB is the one feed that must travel through a proxy. The other two are unaffected and their results remain valid.
SSVC shows INSUFFICIENT DATA	Intel is missing or older than 24 hours. Re-click LOAD EXPLOIT INTEL. The tool withholds an outcome rather than scoring you on a stale picture.
No BOD tier is shown	Automatable or Technical Impact is unknown. Click LOAD CISA VULNRICHMENT, or set Technical Impact yourself in the detail panel.
A tier shows two values with a "?"	Exposure is not set. Set the estate default in the BOD bar or answer per-CVE. Until then the worse case is displayed.

APPENDIX A — THE BOD 26-04 DECISION TABLE

All sixteen rows as implemented, extracted directly from the tool so this table cannot drift from the code. Each row is covered by a unit test.

In KEV	Publicly Exposed	Automatable	Technical Impact	Required timeline
NO	NO	NO	Partial	FIX ON UPGRADE
NO	NO	NO	Total	FIX ON UPGRADE
NO	NO	YES	Partial	60 DAYS
NO	NO	YES	Total	60 DAYS
NO	YES	NO	Partial	60 DAYS
NO	YES	NO	Total	14 DAYS
NO	YES	YES	Partial	14 DAYS
NO	YES	YES	Total	3 DAYS
YES	NO	NO	Partial	14 DAYS
YES	NO	NO	Total	14 DAYS
YES	NO	YES	Partial	14 DAYS
YES	NO	YES	Total	3 DAYS + TRIAGE
YES	YES	NO	Partial	14 DAYS
YES	YES	NO	Total	3 DAYS + TRIAGE
YES	YES	YES	Partial	3 DAYS
YES	YES	YES	Total	3 DAYS + TRIAGE

VERIFY AGAINST THE DIRECTIVE BEFORE OFFICIAL USE This table was transcribed from two independent secondary reproductions (CERT/CC and Nucleus), which agree on all sixteen rows. It has not been diffed against the text on cisa.gov. One widely-circulated vendor summary of this directive is known to be wrong — it states that KEV plus total impact is always a 3-day-with-triage item regardless of exposure, where both table sources put KEV / not exposed / not automatable / total at 14 days. Confirm against the directive itself before this drives a reportable deadline.

APPENDIX B — WHY THREE VERDICTS INSTEAD OF ONE

A single blended score would be easier to read and less honest. The three verdicts answer different questions, and collapsing them would hide which input is actually driving the answer.

Verdict	Optimised for	Breaks down when
Priority P1-P4	Speed. Can an attacker fire this today?	Exploit intel is unloaded or stale — capped at P3 and labelled
SSVC	Defensible decisions. Reproduces CISA's published tree.	Mission & Well-being is unset — shows a range instead of guessing
BOD 26-04	Compliance. What deadline does the directive impose?	Exposure or impact is unknown — shows a span, or withholds the tier

When two of them disagree, look at which input differs. A P1 that lands at 14 days under BOD usually means the asset is not publicly exposed. An ACT under SSVC with no BOD tier usually means Technical Impact is unknown.

Honest framing for a brief

If you are citing this tool in a report, the accurate description is:

"CISA SSVC decision tree with two decision points derived from public exploit intelligence rather than analyst evaluation."

Not "CISA SSVC scoring". The tree is CISA's and is transcribed verbatim; two of the four inputs are derived unless you override them, and the tool labels which is which on every row.

Known limitations, stated plainly

- P2 cannot currently be assigned — the proof-of-concept feed is not wired (section 5).
- Exploitation's PoC definition is narrower than CISA's, which also counts published exploit descriptions and obviously-exploitable vulnerability classes.
- Technical Impact falls back to the CVSS vector as a proxy. CISA does not define it via CVSS at all — which is exactly why the system-owner override exists.
- Absence of a public exploit is absence of evidence. Private and nation-state tooling is not observable from any public feed.

GHOST DEFENSE

Grounded Handling Of Sourced Threat-intel

Zero-install, browser-based defensive security tooling
for resource-constrained organizations.

[ghost2210 . pages . dev](https://ghost2210.pages.dev)

DEFEND ALWAYS.