A Practical Internet-Reachability Test for FedRAMP

Matthew Venne
Chief Technology Officer, stackArmor

July 2026

A plain-language companion to "Internet Reachability at Scale under the FedRAMP VDR/VER Rules."

The short version

FedRAMP asks providers to decide, for every vulnerability finding, whether a payload from the public internet might trigger it. That is broader than asking whether the affected server has a public IP. A crafted request can pass through an application and reach a private database, parser, queue worker, or logging library.

The definition is technically sound. The operational problem is that ordinary scanner findings usually do not say whether a CVE can be triggered after its input crosses a component boundary. At scale, a provider cannot manually research that question for every finding on every asset.

The paper recommends a current triage rule:

THE WHOLE MODEL

A finding is internet-reachable when its vulnerable network service can be triggered directly, or when a known indirect trigger reaches the component and current assessment and continuous-monitoring evidence does not show that the relevant trigger class is prevented before vulnerable processing.

That rule is intentionally practical. It uses information providers already have, promotes exceptional indirect triggers once per CVE, and keeps a more detailed application-wide data-flow model optional.

Why the scanner cannot finish the job

A scanner can usually identify the affected package, CVE, version, and CVSS vector. It may also report the listening service and a short vulnerability description. That is enough to automate much of the direct-exposure question.

The missing fact is the trigger path. A local parser flaw may be completely unrelated to internet input, or it may process every uploaded document. A logging library may have no listener at all but still receive attacker-controlled strings from a public API. CVSS attack vector helps describe the vulnerable component; it does not reliably say whether another component can carry the trigger to it.

This is primarily an upstream metadata problem. Vendors, CNAs, vulnerability publishers, and scanner feeds should distribute a reusable field such as `indirect_internet_trigger`, with provenance and enough trigger detail to support review. Today they generally do not.

Three checks, in order

The operational method asks three questions.

1. **Can the internet directly deliver what this flaw requires?** Match public routing, live listeners, protocol, version, and relevant edge controls to the vulnerability's direct trigger. A public host does not make every installed package reachable.

2. **Is this a known indirect-trigger vulnerability?** Record a reusable CVE-level status: confirmed, ruled out, or unknown. A vendor advisory, threat-intelligence report, exploit analysis, reviewed agentic proposal, or analyst decision can confirm the status once and reuse it across the fleet.

3. **Does the component receive internet-derived content, and is the trigger currently prevented?** A coarse component archetype can answer the first half. Current assessment, DAST, remediation, and retest evidence can support the second half for the trigger or input-control classes they actually cover.

The paper calls that last prevention assertion P_0. If P_0 is current and in scope, a known downstream trigger can be classified NIRV. If the evidence is missing, stale, contradicted, or unrelated to the trigger, P_0 is zero and the finding remains IRV.

The evidence already exists, but it has limits

For Rev5 services, the FedRAMP Continuous Monitoring Playbook requires recurring web-application scanning, coverage of web interfaces or an approved sample, individual tracking of findings, and assessor-validated scanner configuration. Providers maintain that evidence and remediate findings. Assessors review its scope and integrity and verify provider-run scan results.

That division of labor matters. The provider produces and refreshes the evidence. The 3PAO collects or reviews it during assessment and validates the method. The model does not ask the assessor to invent a complete application data-flow graph or research every CVE on the provider's behalf.

A current, properly scoped DAST program is meaningful sanitation and validation evidence. It is also not universal proof. Automated testing can miss authenticated paths, unusual formats, new trigger techniques, and code that the test plan never exercises. The OWASP Web Security Testing Guide makes the same point: automated scanning belongs in a balanced testing program.

The practical line is straightforward:

- A clean, current scan with validated scope, relevant coverage, remediated findings, and successful retest can support P_0 for the tested class.
- A generic clean scan with unknown scope cannot.
- A new relevant finding, failed retest, material application change, uncovered interface, alternate path, or new trigger technique revokes or narrows P_0.

Why reachability converges toward accessibility

The model makes two explicit triage presumptions.

First, a CVE whose indirect-trigger status is unknown is not automatically promoted across every component that ever handles internet-derived data. Doing that would make nearly every connected system IRV and erase the prioritization the field is meant to provide. The unknown status remains visible and is queued for risk-ranked enrichment.

Second, a current P_0 assertion is trusted within its validated scope until contrary evidence or a material change revokes it. This reuses the assessment and ConMon work providers already perform.

The resulting operational IRV list therefore converges toward the directly internet-accessible surface, plus known indirect triggers that are not covered by current prevention evidence. Reachability and accessibility do not mean the same thing. They produce similar operational lists because tested prevention closes many downstream trigger paths, while the exceptional paths are promoted explicitly.

That is a conditional convergence, not a promise that downstream content is always safe. Its residual risks are equally explicit: missing upstream trigger metadata and incomplete or stale prevention evidence.

Three examples

- **SQL construction behind an application tier.** Internet-derived values reach a private database. The validated DAST program covers SQL injection on the relevant interface, relevant findings were remediated and retested, and no bypass or material change is known. The current P_0 assertion makes the database finding NIRV. A new SQL-injection finding or uncovered interface resets P_0, and the finding returns to IRV.

- **A logging-expression trigger.** Threat intelligence confirms that a crafted string can cross process boundaries and trigger the affected logging library. The component receives internet-derived log content. Ordinary SQL-injection and cross-site-scripting DAST does not cover this mechanism, so P_0 is zero and the finding is IRV.

- **An administrative service behind IAP or VPN.** A separate, bypass-free edge authorizes the connection before traffic reaches the backend and is the backend's only ingress path. That can remove the backend's direct internet path. A login implemented by the backend application does not do the same thing because the request has already reached the component; application authentication instead informs exploitability or an attested mitigation.

What happens when a new trigger appears

Log4Shell is the useful pattern. Before the trigger was understood, ordinary scanner output and web testing did not say that a string written to a log could activate a private library. Once the mechanism became known, the CVE-level indirect-trigger status could be promoted across every affected deployment archetype that received attacker-controlled log content.

That promotion is why trigger metadata belongs upstream. Research it once, publish it with provenance, and let every scanner and vulnerability-management platform reuse it. Agentic analysis can propose missing profiles, but it should cite sources, preserve unknowns, and require stronger

review before concluding that no indirect trigger exists. An agent cannot recover a technical fact that no authoritative source disclosed.

When new intelligence shows that an existing prevention assertion does not cover the trigger, P_0 is revoked and the affected findings are re-evaluated. The control is not allowed to remain true merely because last month's scan was clean.

The optional future state is not today's requirement

A richer system could import or derive a deployment graph, record how content moves between components, attach transformations such as parameterized queries to each path, and bind tests to the deployed artifact. That would improve precision and make stronger NIRV claims portable between tools.

The industry does not yet have a broadly adopted way to combine those artifacts into an independently verifiable, vulnerability-specific reachability record that vulnerability-management platforms can consume. The full paper includes a non-normative equation showing how such a system could work. It is an interoperability direction, not a prerequisite for the current method and not a claim that providers must build a complete application-wide graph now.

A brief word about WAFs

FedRAMP allows a verified perimeter control that blocks the triggering payload before vulnerable processing to remove IRV status. The paper recommends a more durable operational posture for volatile WAF rules: keep the intrinsic reachability visible and attach a signed mitigation or VEX assertion describing the rule, its scope, and its evidence. If the rule is bypassed, disabled, or routed around, the record fails visibly instead of leaving behind a stale NIRV classification. KEV remediation dates continue to apply even when a vulnerability is fully mitigated.

The bottom line

1. Reachability is evaluated per finding, not assigned once to an asset.
2. Direct exposure can be automated from network, listener, protocol, and runtime evidence.
3. Exceptional indirect triggers should be enriched once per CVE and reused across the fleet; today's scanners rarely provide that metadata.
4. Existing assessment, DAST, remediation, and retest evidence supports a scoped, revocable prevention assertion without requiring a complete data-flow graph.
5. Under those declared presumptions, the operational IRV list converges toward direct accessibility plus known uncovered indirect triggers.

The method is not exhaustive exploit-path proof. It is a scalable triage model that states exactly what it assumes, preserves the unknowns, and identifies the upstream metadata improvement that would make the result more precise.

Read the full technical paper.