
White Paper | September 2026

Coalfire's stance on Potential Agency Impact (PAIN) ratings

A methodology white paper for **FedRAMP 20x** and the Consolidated Rules for 2026 (CR26)

AUTHOR

Michael Allen | Director – Security Operations

Contents

Executive summary	03
Key positions	04
Key terms and how the evaluation works	05
Regulatory basis	07
Evaluation factors and related rules	08
Regulatory definitions of customer effect	09
The interpretive gap	09
Coalfire's stance: the N3 default posture	10
Two dimensions, evaluated separately	11
Coalfire's customer-effect calibration	12
Documentation standard	13
Vulnerability PAIN vs. incident PAIN	14
Recommended remediation timeframes at Class C	15
Worked examples	16
Classifying assets ahead of time	18
Common pitfalls	19
Summary	19
Appendix: sources	20

Executive summary

SCOPE ASSUMPTION: FEDRAMP CERTIFICATION CLASS C

This paper assumes the customer's cloud service offering holds a Class C Certification. Force levels and timeframes vary by Certification Class; a provider under Class A, B, or D should re-check the class-specific figures cited here before using this paper as a compliance basis.

FedRAMP's Consolidated Rules for 2026 (CR26) require providers subject to the applicable VER rules to rate the potential agency impact of each detected vulnerability on a five-level scale, N1 through N5. The rule, VER-EVA-EPA, defines what each rating means. A companion rule, VER-EVA-EFA, states that providers SHOULD consider a defined set of contextual factors, including criticality, when making that determination. Neither rule prescribes a quantitative scoring model for combining those factors into a rating.

Coalfire fills that space with a documented, repeatable, and auditable methodology built on two inputs intended to be distinct: asset criticality and Common Vulnerability Scoring System (CVSS) Environmental severity. This pairing is Coalfire's implementation choice — CR26 requires contextual evaluation, not CVSS Environmental scoring specifically, and other defensible methods may exist. Coalfire's default posture starts every new finding at N3, a mid-scale rating, until the evaluation produces evidence for a different rating.

This paper lays out the regulatory basis for that position, the calibration Coalfire proposes, the documentation standard behind it, and how vulnerability-level PAIN relates to, but is evaluated separately from, incident-level PAIN under FedRAMP's incident rules.



Coalfire's default posture starts every new finding at N3, a mid-scale rating.

Key positions

PAIN runs from N1 (lowest potential impact) to N5 (highest); the positions below use that orientation throughout.



Key terms and how the evaluation works

Terms marked Coalfire are Coalfire methodology terms, not FedRAMP-defined terms.

TERM	PLAIN-LANGUAGE MEANING
PAIN	FedRAMP's five-level rating of impact. N1 is lowest, N5 is highest. Defined in VER-EVA-EPA.
Customer effect	How badly an affected agency would be hurt: minimal, narrow, disruptive, or debilitating.
Agency scope	Whether the effect reaches one agency or more than one.
IRV / NIRV	Internet-reachable vulnerability or not — reachable even indirectly through other components.
LEV / NLEV	Likely exploitable vulnerability or not, judged for the specific cloud service offering.
Asset criticality COALFIRE	How important an asset is: Low, Moderate, or High. FedRAMP lists this as a factor but does not define the levels.
CVSS Base score	A vulnerability's severity in the abstract, before considering where it runs.
CVSS Environmental score	The same vulnerability's severity on a specific asset, adjusted for data sensitivity and configuration.
Calibration table COALFIRE	Coalfire's table that turns criticality and Environmental severity into a customer effect (Figure 3).
Accepted vulnerability	Not, or will not be, fully mitigated or remediated within 192 days of evaluation.
FedRAMP Reportable Incident	Affects, or is likely to affect, confidentiality or integrity of federal customer data.
Bridge rules	VER-TFR-IRI and VER-TFR-NRI — say when a vulnerability should or may be treated as an incident.
MUST, SHOULD, MAY	FedRAMP force levels: mandatory, expected unless justified otherwise, and optional (IETF RFC 2119).
Certification Class	FedRAMP's Class A–D tiers. Timeframes and force levels vary by class; this paper assumes Class C.

How the evaluation works

Each detected vulnerability under CR26 is evaluated for three things: internet reachability, likely exploitability, and potential agency impact (PAIN).




Those three results together set the recommended remediation timeframe and feed required reporting.

In specific combinations, the bridge rules say the vulnerability should or may also be treated as a FedRAMP Reportable Incident. At that point a separate incident evaluation begins, with its own PAIN rating and reporting timeframes.

See Figure 1, next page →

Regulatory basis

FIGURE 1 How a Class C vulnerability moves through CR26 evaluation



Shaded boxes are incident (IEC) rules; the dashed box is the VER bridge check.

VER-EVA-EPA, “Estimate Potential Agency Impact,” is a MUST-level rule under CR26, applicable across 20x and Rev5 Certification Types, with a required obtain-and-maintain date of December 7, 2026 (CISA BOD 26-04). The published grace period ends March 7, 2027.

VER-EVA-EPA, VERBATIM

- N1:** Minimal customer effect on one or more agencies.
- N2:** Narrow customer effect on one or more agencies.
- N3:** Disruptive customer effect on one agency.
- N4:** Debilitating effect on one agency OR disruptive effect on more than one agency.
- N5:** Debilitating customer effect on more than one agency.

Source: FedRAMP CR26, Vulnerability Evaluation and Reporting, VER-EVA-EPA. Official launch 2026-06-24. Retrieved September 25, 2026.

Evaluation factors and related rules

VER-EVA-EFA (SHOULD-level): providers SHOULD consider at least these factors when evaluating detected vulnerabilities in context.

1 Criticality. How important are the systems or information that might be impacted?

2 Reachability. How might a threat actor reach the vulnerability, and how likely is that?

3 Exploitability. How easy is it to exploit, and how likely is that?

4 Detectability. How easy is it for a threat actor to become aware of it?

5 Prevalence. How much of the cloud service offering is affected?

6 Privilege. How much privileged access is granted or gained by exploiting it?

7 Proximate vulnerabilities. How does it interact with previously detected, partially or fully mitigated vulnerabilities?

8 Known threats. How might already known threats leverage the vulnerability?

RELATED RULES (CLASS C)

RULE	WHAT IT GOVERNS (CLASS C)
VER-TFR-EVU	Class C providers SHOULD evaluate all vulnerabilities within 5 days of detection.
VER-TFR-IRI	SHOULD treat internet-reachable, likely exploitable findings rated above N3 as a Reportable Incident until mitigated to N3 or below.
VER-TFR-NRI	MAY treat non-internet-reachable, likely exploitable findings rated N5 as a Reportable Incident until mitigated to N4 or below.
VER-RPT-VDT	MUST include historical and current PAIN, each reduction, and the next planned reduction for every non-accepted vulnerability.
VER-TFR-MRH	Class C providers SHOULD publish historical vulnerability activity in JSON at least every 14 days.
VER-TFR-MAV	MUST categorize a vulnerability as accepted if not mitigated or remediated within 192 days of evaluation.

Regulatory definitions of customer effect

CUSTOMER EFFECT	FEDRAMP DEFINITION (SUMMARIZED)
Minimal	Noticeable by some users only; includes minor inconveniences such as reduced performance.
Narrow	Interrupts use for some users for less than 12 hours, or compromises an extremely limited amount and type of federal customer data.
Disruptive	Interrupts use for many users for less than 24 hours, or compromises large amounts or many types of federal customer data.
Debilitating	Interrupts use for most users, or compromises the integrity or confidentiality of most federal customer data. If the adverse effect is unknown, it should be treated as debilitating until proven otherwise.

That last line supports treating an unknown effect conservatively; it does not itself dictate where a provider’s default should sit on the five-level scale. FedRAMP’s debilitating definition states an unknown adverse customer effect should be treated as debilitating until proven otherwise — but it does not expressly require an unclassified asset or missing CVSS Environmental score to automatically receive a PAIN-5 rating, nor does it foreclose a mid-scale starting point pending evaluation.

Coalfire’s N3-default posture (Section 3) is Coalfire’s own methodology: it starts a new finding at Disruptive-equivalent impact and moves the rating up or down as evidence comes in. FedRAMP neither mandates nor precludes that choice of starting point.

The interpretive gap

VER-EVA-EPA is precise about outcomes; VER-EVA-EFA adds contextual factors. Neither prescribes a quantitative scoring model — a formula, matrix, or weighting scheme — for turning those factors into one of the five N-ratings for a specific finding on a specific asset. That interpretive space is where a provider’s methodology becomes either a defensible, auditable practice, or a source of inconsistent, ungrounded ratings.

The N3 default posture

DEFAULT POSTURE

Coalfire's methodology assigns a new finding N3, a mid-scale rating equivalent to a Disruptive customer effect on a single agency — not N1 and not N5 — until an evaluation produces evidence for a different rating.

This applies with equal force when an asset has not yet been assigned a criticality classification, or when a CVSS Environmental score has not yet been computed: missing data holds the rating at N3 rather than moving it in either direction.

Moves down toward N1

with documented evidence of compensating controls, reduced environmental severity, or narrow scope.

Moves up toward N5

with documented evidence of internet reachability, likely exploitability, high asset criticality, or elevated environmental severity.

This is Coalfire's methodology choice. FedRAMP does not mandate this specific default; Coalfire applies it as a documented, evidence-driven starting point rather than an assumed worst case or best case — treating an unrated finding as either low-impact or worst-case by omission, rather than by evaluated evidence, is the failure mode this default is built to prevent.

Two dimensions, evaluated separately

VER-EVA-EPA's rating definitions combine two separate dimensions, customer effect and agency scope. Evaluate each on its own, then read the resulting N-rating from Figure 2.

FIGURE 2 Customer effect ✕ agency scope → PAIN rating

CUSTOMER EFFECT	ONE AGENCY	MORE THAN ONE AGENCY
Minimal	N1	N1
Narrow	N2	N2
Disruptive	N3	N4
Debilitating	N4	N5

This combination is FedRAMP rule text. Darker cells are higher customer effect.

Scope only changes the outcome at the Disruptive and Debilitating tiers. FedRAMP's own rule text for N1 and N2 reads "one or more agencies" without a split by scope, so applying a scope test to a Minimal or Narrow finding does not change the answer and is unnecessary evaluation effort.

Coalfire’s customer-effect calibration

Coalfire’s proposed method for meeting VER-EVA-EPA’s requirement to evaluate in context. It is not FedRAMP rule text; the cell values are Coalfire’s calibration. It crosses asset criticality against the CVSS Environmental score for the specific finding on the specific asset — not the abstract CVSS Base score.

FIGURE 3 Asset criticality × CVSS Environmental severity

ASSET CRITICALITY	CVSS ENVIRONMENTAL SEVERITY			
	Low 0.1 to 3.9	Medium 4.0 to 6.9	High 7.0 to 8.9	Critical 9.0 to 10.0
High	Narrow	Disruptive	Disruptive	Debilitating
Moderate	Minimal	Narrow	Disruptive	Disruptive
Low	Minimal	Minimal	Narrow	Narrow

Coalfire methodology, not FedRAMP rule text. Darker cells are higher customer effect.

Requiring documented non-overlap. CVSS Environmental scores already incorporate asset-specific CIA requirements. For the calibration to be sound, asset criticality must be documented on a distinct basis — mission dependency, customer-facing status, federal data classification tier — rather than assumed independent by default. If asset criticality instead restates the CIA requirements already folded into the Environmental score, the calibration weights one signal twice and inflates the rating without a second, distinct basis for doing so.

Documentation standard

Every PAIN determination should be recorded with enough detail to reconstruct the rating during an assessment.

FIELD	WHAT IT CAPTURES
Asset criticality	Low, Moderate, or High, and the basis for that classification
CVSS Environmental score and vector	The score used to derive customer effect, not the Base score
Customer effect category	Minimal, Narrow, Disruptive, or Debilitating, from the Section 3.3 calibration
Agency scope	Single agency or multi-agency, and which agencies where relevant
N-rating	The final determination, N1 through N5
Default applied	Whether N3 was assigned because criticality or Environmental scoring was not yet available
Evaluator / Timestamp	Who made the determination, and when
Reduction history	Time and N-rating of each completed and evaluated reduction, if it has moved since initial evaluation
Next planned reduction	Estimated time and target N-rating of the next planned reduction
Overdue status	Whether the finding is, or is likely to become, overdue, with an explanation if so

These fields support and explain the required VER-RPT-VDT reporting fields — tracking identifier, detection time, IRV/NIRV and LEV/NLEV determinations, and final disposition — for every non-accepted vulnerability.



Vulnerability PAIN vs. incident PAIN

The vulnerability-level PAIN this paper covers, under VER-EVA-EPA, is a separate determination from incident-level PAIN under Incident Evaluation and Communication (IEC). A high vulnerability PAIN does not, by itself, make a finding a FedRAMP Reportable Incident.

HOW AN INCIDENT GETS ITS OWN PAIN RATING

Under IEC-CSO-EFR, providers MUST promptly evaluate incidents (not vulnerabilities) for effect on confidentiality or integrity of federal customer data. A reportable incident then gets its own PAIN under IEC-CSO-EFI, using the same five customer-effect tiers as VER-EVA-EPA.

IEC-CSO-DPR, VERBATIM

“Providers MUST treat FedRAMP Reportable Incidents as if they have a Potential Agency Impact N-rating (PAIN) of 5 UNLESS they promptly estimate the PAIN rating following the rule in IEC-CSO-EFI.”

Source: FedRAMP CR26, Incident Evaluation and Communication, IEC-CSO-DPR. Official launch 2026-06-24.

At Class C: PAIN-3 to PAIN-5 incidents carry a 1-hour Initial Report, 6-hour Ongoing Reports, and a 6-hour Final Report. PAIN-1 and PAIN-2 incidents carry longer timeframes (24 hours to 1 business day).

The Class C bridge to incident treatment

VER-TFR-IRI (SHOULD)

Applies only when the vulnerability is internet-reachable, likely exploitable, AND rated above N3. All three must hold.

VER-TFR-NRI (MAY)

Applies only when the vulnerability is likely exploitable, NOT internet-reachable, AND rated exactly N5.

Evaluation timing is a SHOULD, not an absolute requirement. VER-TFR-EVU states Class C providers SHOULD evaluate all vulnerabilities, including PAIN, within 5 days of detection — a deviation from that window is not a MUST-level violation unless another rule makes it mandatory.

Recommended remediation timeframes at Class C (VDR-TFR-PVR)

PAIN RATING	LEV + IRV (reachable, exploitable)	LEV + NIRV (not reachable, exploitable)	NLEV (not likely exploitable)
N5	2 days	4 days	16 days
N4	4 days	8 days	64 days
N3	16 days	32 days	128 days
N2	48 days	128 days	192 days


SHOULD-level recommended timeframes at Class C, not MUST-level deadlines. No N1 row — mitigate during routine operations per VDR-TFR-RMN.

Worked examples

Each example follows the evaluation left to right: inputs, customer effect (Figure 3), scope, PAIN (Figure 2), and Class C timeframe (Section 6).


EXAMPLE A High criticality, high severity, multi-agency

A shared database holds federal CUI for three agency customers. A detected CVE computes a High CVSS Environmental score. The calibration gives Disruptive; scope is multi-agency. Per Figure 2, Disruptive plus multi-agency is N4. The Class C timeframe is 4, 8, or 64 days depending on IRV/LEV.



EXAMPLE B Missing classification


A newly onboarded asset has not yet been assigned a criticality classification when a CVE is detected on it. The calibration cannot run without that input, so the finding is rated N3 by default and stays there until the asset is classified or other evidence supports a different rating.



Worked examples (continued)

EXAMPLE C Low criticality, high severity in the abstract

An isolated internal tool with no federal data has a CVE with a CVSS v3.1 Base score of 9.8 (Critical). Lowering only the CIA requirements to Low produces an Environmental score of 8.0 (High). Because the vulnerable component is reachable only locally, requires high privileges, and has limited data impact, the Environmental score with modified base metrics is 2.7 (Low). The calibration gives Minimal; scope does not apply. The rating is N1, handled during routine operations (VDR-TFR-RMN).



EXAMPLE D Reclassification changes the rating

An asset was classified Moderate criticality with a Medium Environmental score → Narrow → N2. Six months later it begins serving a second agency and is upgraded to High criticality. Re-evaluated, the same Medium score now maps to Disruptive, and with multi-agency scope the finding moves to N4. The historical and current PAIN fields record this as a rating increase tied to the reclassification event, not a new finding.

Figure 7a. Before reclassification: N2





Figure 7b. After reclassification: N4



Classifying assets ahead of time

The pre-classification standard. If asset criticality is assigned only after a finding or incident appears, the classification can be shaped, knowingly or not, by the outcome a team wants. Coalfire's standard is to classify every asset before any finding arrives and lock the classification used for each evaluation. This is Coalfire's recommended practice, not a FedRAMP requirement.

■ Make criticality a required inventory attribute, with a documented basis, owner, and assignment date.

■ Publish the Low / Moderate / High criteria before applying them, version-controlled.

■ Snapshot the classification at evaluation time; later changes apply going forward only.

■ Measure coverage — unclassified assets hold findings at N3, so gaps stay visible.

■ Gate onboarding on classification — no asset enters production before it is classified.

■ Map agency scope in advance so scope is a lookup, not an investigation.

■ Control classification changes during an active incident; require documented approval.

■ Review on trigger events and quarterly as a backstop (Coalfire recommendation).



Scaling beyond single-finding triage. Most inputs do not need to be recomputed per finding.

Asset criticality is a property of the asset, not the finding — classify once per asset or asset group.

CVSS Environmental scoring is automatable per asset group once CIA modifiers are set in tooling.

Two inputs resist automation: the initial criticality judgment call, and confirming agency scope.

Both inputs are persistent, not periodic — reclassify on business-change events, not only on schedule.

Common pitfalls

- 1 Using CVSS Base instead of Environmental.** Base ignores the asset entirely and produces the same rating regardless of what it sits on.
- 2 Double-counting criticality.** If criticality restates the CIA requirements already in the Environmental score, the calibration weights one signal twice.
- 3 Applying the scope split to Minimal or Narrow findings.** FedRAMP's own text for those two ratings does not distinguish single- from multi-agency.
- 4 Treating an unclassified asset as low-impact, or assuming worst case, by default.** Coalfire's methodology holds an unrated finding at N3 until evidence, not its absence, moves the rating.
- 5 Letting a rating sit static.** A rating should move when criticality or scope changes, not only when the underlying CVE changes.

SUMMARY

VER-EVA-EPA gives providers a five-level scale, and VER-EVA-EFA states they SHOULD weigh a defined set of contextual factors, but neither rule prescribes a quantitative scoring model. Coalfire's stance fills that space with two distinct inputs (asset criticality and CVSS Environmental severity), an N3 default that moves only with evidence, and a documentation standard that reconstructs every rating without relying on institutional memory. Vulnerability PAIN and incident PAIN remain separate determinations, tied together only at the bridge rules in Section 5. The Class C remediation timeframes are SHOULD-level, not MUST-level deadlines — document deviations rather than treating the table as a hard stop.

Appendix: sources

All rules retrieved September 25, 2026. “Latest changelog” is the most recent entry shown for the rule on that date.

RULE VERSIONS USED

RULE OR TERM	RULESET	LATEST CHANGelog
VER-EVA-EPA, -EFA, -ELX, -EIR	Vulnerability Evaluation and Reporting	2026-06-24, official launch
VER-RPT-VDT	Vulnerability Evaluation and Reporting	2026-06-24, official launch
VER-TFR-EVU, -IRI, -NRI, -MAV, -MRH	Vulnerability Evaluation and Reporting	2026-06-24, official launch
VDR-TFR-PVR	Vulnerability Detection and Response	2026-07-01, terms/related controls added; requirement unchanged
VDR-TFR-RMN, VDR-CSO-RES	Vulnerability Detection and Response	2026-06-24, official launch
IEC-CSO-EFR	Incident Evaluation and Communication	2026-07-02, terminology changed Response → Communication
IEC-CSO-EFI, -DPR, -IIR, -OIR, -FIR	Incident Evaluation and Communication	2026-06-24, official launch
Debilitating / Disruptive / Narrow / Minimal	FedRAMP Definitions	2026-06-24, official launch

SOURCE DOCUMENTS

FedRAMP Consolidated Rules for 2026, Vulnerability Evaluation and Reporting — fedramp.gov/2026/reference/vulnerability-evaluation-and-reporting

FedRAMP Consolidated Rules for 2026, Vulnerability Detection and Response — fedramp.gov/2026/reference/vulnerability-detection-and-response

FedRAMP Consolidated Rules for 2026, Incident Evaluation and Communication — fedramp.gov/2026/reference/incident-evaluation-and-communication

FedRAMP Definitions (Customer Effect terms, Potential Agency Impact) — fedramp.gov/2026/definitions

CVSS v3.1 Specification, Environmental Metric Group — first.org/cvss/v3-1/specification-document

CISA Binding Operational Directive 26-04 — cisa.gov/news-events/directives/bod-26-04-prioritizing-security-updates-based-risk

This paper reflects Coalfire's interpretive methodology under the current FedRAMP Consolidated Rules for 2026 and is not a substitute for a provider's own legal or compliance review.



Ready to talk through your PAIN methodology?

Coalfire's advisory team can review your asset criticality classifications, CVSS Environmental scoring, and documentation against what a FedRAMP assessor will expect to see.

Talk to our team today.