Before the PAIN Equation: Deriving Security-Requirements Ceilings from Intended Federal Information Types

Matthew Venne
Chief Technology Officer, stackArmor

July 2026

Abstract

Potential Agency Impact (PAIN) calculations require a Confidentiality, Integrity, and Availability Security Requirements vector for the affected asset. This paper calls the uncapped vector the asset's *security-impact profile*. A CSP may derive it through a role-based archetype catalog, direct per-objective assessment, or another governed method, but no architectural assignment method can by itself answer a different question: which federal consequences are possible in a particular agency's intended use of a Cloud Service Offering (CSO)?

This paper specifies the missing upstream derivation. It begins with the federal information types the CSO is designed and authorized to process and the information types an agency intends and is authorized to place in that offering. Confirmed overlap is mapped to the provisional objective values and special factors in NIST SP 800-60. The applicable values are aggregated independently for Confidentiality, Integrity, and Availability, preserving the full dimensional vector required by FIPS 199. That vector becomes a Security Requirements ceiling. The effective requirement for an asset is the per-objective minimum of its system-relative security-impact profile and the intended-use ceiling.

The separation matters. Security-impact profiles remain reusable across agency deployments because they describe the component inside the CSO, not every possible customer data type. A provider that uses archetypes can reuse those role mappings; a provider that assigns the three objectives directly can reuse its governed assessment. The ceiling supplies the agency-specific consequence context without turning a system's overall FIPS 199 impact label into an all-equal vector. Generic free text, file upload, or integration capability does not authorize hypothetical information types. Multi-agency assets aggregate only the definite agencies they can affect. The assessment measures adverse effect on federal operations, federal assets, and individuals; a CSP's separate brand or corporate risk does not silently raise agency impact.

This paper ends at the effective Security Requirements vector. It does not choose a PAIN aggregation formula, word threshold, or N-rating. Those downstream policy choices belong to the PAIN method and its calibration. It also separates that consequence vector from Certification Class: Class describes the operational assurance and response commitments the provider supplies, while the vector describes what the agency has at stake. Where those signals diverge, the agency and CSP must resolve the difference through authorization scope, service design, contract terms, or explicit risk acceptance rather than altering either signal. For an Availability objective rated High, the paper also recommends a specific availability-assurance case that translates the agency consequence into service design, disaster-recovery, and contractual evidence without prescribing one topology or uptime percentage.

Contents

1 Status, scope, and relationship to PAIN

This document is informational. It does not define a FedRAMP requirement, replace an agency categorization, or make NIST SP 800-60 recommendations mandatory. It proposes a deterministic bridge from federal information categorization to the Security Requirements inputs consumed by the PAIN method.

The key words MUST, MUST NOT, SHOULD, SHOULD NOT, and MAY describe this proposed method as in RFC 2119. They do not create obligations under FedRAMP, FIPS 199, or NIST guidance.

Cloud Service Offering (CSO)
The bounded cloud service whose intended federal use is being assessed.

Information type
A category of federal information identified under NIST SP 800-60, with provisional Confidentiality, Integrity, and Availability impact recommendations.

Applied profile
The objective vector retained after an agency confirms that an information type applies and evaluates its special factors and actual use.

Asset security-impact profile
The uncapped, independently assignable CR/IR/AR vector for an affected asset. It may be derived from an archetype, direct per-objective assessment, or another governed method.

Security-requirements ceiling
The greatest C/I/A consequence supported by the confirmed intended information use in the asset's agency scope.

Effective requirements
The per-objective minimum of the asset security-impact profile and the ceiling. This is the CR/IR/AR vector handed to PAIN.

SCOPE BOUNDARY

The paper stops at effective CR/IR/AR. Whether a particular effective vector can produce Minimal, Narrow, Disruptive, or Debilitating under a selected PAIN formula is a property of that downstream formula and its governed thresholds. It is not a property established by NIST SP 800-60 or FIPS 199 alone.

2 The missing first step

FedRAMP requires providers to evaluate vulnerabilities in the context of the CSO and estimate potential impact on government customers. CVSS Environmental Security Requirements supply a standard vocabulary for stating how strongly the affected environment requires confidentiality, integrity, and availability. Neither source, however, automatically provides the three values for a particular asset in a particular agency deployment.

Why this step still matters under FedRAMP CR 2026. The derivation in this paper deliberately follows the substance of the FIPS 199 and NIST SP 800-60 categorization process: identify the applicable information types, determine potential impact separately for Confidentiality, Integrity, and Availability, retain the dimensional vector, and separately derive the overall high-water-mark category.

FedRAMP's Consolidated Rules for 2026 (CR 2026) replace the historical provider Low/Moderate/High Impact Level designation with Certification Classes A–D and expressly state that Class and Impact Level do not directly correlate. In practical terms, this removes the old shortcut by which a provider's certification label appeared to carry a FIPS 199 category through the provider certification process. It does not remove the agency's separate responsibility to categorize its federal information system [FEDRAMP-CLASS, FEDRAMP-2026].

That decoupling makes the categorization reasoning no less important for PAIN. PAIN still asks the CSP to estimate the potential agency effect of exploiting a vulnerability, and CVSS Environmental scoring still needs objective-specific CR/IR/AR values. A Certification Class alone cannot supply those values. This paper preserves the FIPS 199 consequence analysis needed to derive them without turning the Class into an agency system category or the CSO certification into an agency authorization.

When we originally wrote the PAIN method, we recommended asset archetypes: governed role descriptions such as system-of-record database, identity service, orchestrator, message broker, or disposable cache. We did so to replace unexplained asset-owner declarations with reusable classifications whose rationale could be reviewed and audited. In that original formulation, however, we asked the archetype's CR/IR/AR vector to carry two different kinds of information:

1. **Architectural consequence:** what role does this component perform inside the CSO, and how much can its compromise affect the offering?

2. **Federal mission consequence:** what information and government use are actually in scope for the affected agency or agencies?

That coupling was a pragmatic first step, but its limitation becomes visible when the same CSO and architectural roles are reused across agencies. Combining both questions inside every archetype makes the catalog customer-specific. A database archetype may become High for one agency, Moderate for another, and then change again when the CSO adds a customer. The resulting catalog is difficult to reuse and difficult to audit.

The refinement proposed here preserves role-based archetypes as one useful mechanism while generalizing the first answer into an asset security-impact profile. The profile holds the architectural answer regardless of how it was derived. A Security Requirements ceiling derived from intended federal information use holds the second. PAIN receives their per-objective intersection.

IN PLAIN TERMS

A primary database can be architecturally important even when the current federal use is only Moderate. Keep the database role stable. Derive the agency's three-dimensional ceiling separately. The database keeps its identity, while its effective PAIN requirements reflect what is actually at stake for that agency.

3 Risk has a point of view

FIPS 199 defines impact through adverse effects on organizational operations, organizational assets, and individuals. NIST SP 800-60 applies that federal frame to information types and agency missions. FedRAMP's PAIN vocabulary likewise asks about effects on agencies using the cloud service.

Risk is therefore not an intrinsic property of a byte sequence. The same information can carry different consequences for different stakeholders:

Stakeholder	**Question**	**Treatment in this method**
Federal agency	Would disclosure, corruption, or loss cause limited, serious, or severe or catastrophic adverse effect on the mission, federal assets, or individuals?	This is the impact measured by the ceiling and handed to PAIN.
CSP	Would the same event harm brand, market position, revenue, contractual posture, or corporate reputation?	Important enterprise risk, but not automatically federal agency impact. Include it only where it creates a documented federal consequence.
Other customer	Would the event harm a non-federal customer's operations or obligations?	Outside federal PAIN unless that consequence also affects an agency in the assessed scope.

Table 1: Stakeholder consequences are related but not interchangeable.

A Low federal Confidentiality objective does not mean disclosure is desirable or costless. It means the agency has determined that unauthorized disclosure would have only limited adverse effect under the FIPS 199 frame. A CSP may rationally assign greater internal importance to avoiding the same disclosure. That separate judgment belongs in the CSP's enterprise risk process rather than being smuggled into PAIN as if the government made it.

4 NIST SP 800-60 as provisional evidence

NIST SP 800-60 Volume II Revision 1 contains a catalog of management-and-support and mission-based information types. Each scored record supplies a provisional C/I/A profile, objective-level rationale, special factors, and a recommendation. The values are starting points. The agency identifies applicable information types, evaluates actual use, applies special factors, considers aggregation and system context, and documents adjustments.

4.1 What the catalog does and does not decide

For an information type t, write its provisional profile as

$$p_t = (p_t^C, p_t^I, p_t^A), \qquad p_t^o \in \{L, M, H\}.$$

After the agency evaluates actual use and special factors, the applied profile is

$$q_{a,t} = (q_{a,t}^C, q_{a,t}^I, q_{a,t}^A).$$

The method MUST preserve both values and explain every difference. A provisional High is not automatic proof that the agency's use is High. A provisional Low does not prevent a documented special factor or aggregate from raising the applied objective. A candidate information type does not participate in the math until its applicability is confirmed.

The catalog does not cover classified information or national-security systems. Those uses remain outside this method.

4.2 High-impact paths in the catalog

A source-pinned conversion of Volume II contains 170 records: 168 scored information types and two delivery mechanisms without standalone profiles. Eleven of the 168 scored types begin with at least one provisional High objective:

ID	Information type	Provisional vector
D.2.2	Key Asset and Critical Infrastructure Protection	C:H/I:H/A:H
D.2.3	Catastrophic Defense	C:H/I:H/A:H
D.2.4	EOP Executive Functions	C:H/I:M/A:H
D.3	Intelligence Operations (Unclassified Domestic Intelligence)	C:H/I:H/A:H
D.4.1	Disaster Monitoring and Prediction	C:L/I:H/A:H
D.4.4	Emergency Response	C:L/I:H/A:H
D.5.1	Foreign Affairs	C:H/I:H/A:M
D.5.3	Global Trade	C:H/I:H/A:H
D.11.4	Space Operations	C:L/I:H/A:H
D.14.4	Health Care Delivery Services	C:L/I:H/A:L
D.17.2	Legal Defense	C:M/I:H/A:L

Table 2: The eleven information types with a provisional High objective.

Only four provisionally begin at H/H/H. Another 94 scored records contain objective-specific special-factor language that explicitly permits High under stated conditions. Sixty-three have no explicit High path in either their provisional profile or special factors.

These counts require careful interpretation. Provisional High is uncommon. Conditional High is available across a broader part of the catalog, but a conditional path does not make the information type High by itself. Common conditions concern life safety, primary mission failure, critical infrastructure, time-critical emergency or operational decisions, catastrophic market disruption, and disclosure or corruption that enables catastrophic harm. Each requires an affirmative fact about the agency's intended use.

SCOPE BOUNDARY

The catalog supports "provisional High is rare" and "conditional High requires specific facts." It does not support the claim that only a small fraction of information types have any conceivable High path: 105 of 168 have either a provisional or explicit conditional path. Nor does the catalog establish how many agencies actually meet those conditions.

5 Preserving the FIPS 199 dimensions

5.1 Information-type profiles

FIPS 199 expresses the security category of information as a tuple:

$$\mathrm{SC}_t = \{(C, q_t^C), (I, q_t^I), (A, q_t^A)\}.$$

The dimensions are not interchangeable. An information type may be High for Integrity and Low for Confidentiality and Availability. That asymmetry is evidence, not noise.

5.2 System aggregation

For the set of applicable information types T in a system, FIPS 199 takes the high-water mark independently for each objective:

$$\mathrm{SC}_{\mathrm{system}}(C) = \max_{t \in T} q_t^C, \tag{1}$$

$$\mathrm{SC}_{\mathrm{system}}(I) = \max_{t \in T} q_t^I, \tag{2}$$

$$\mathrm{SC}_{\mathrm{system}}(A) = \max_{t \in T} q_t^A. \tag{3}$$

NIST SP 800-60 then calls for review of aggregation, critical system functionality, and other system factors, with documented adjustment where warranted.

The system's overall impact level is the highest of the three resulting objectives:

$$\mathrm{Impact}_{\mathrm{system}} = \max\{\mathrm{SC}_{\mathrm{system}}(C), \mathrm{SC}_{\mathrm{system}}(I), \mathrm{SC}_{\mathrm{system}}(A)\}.$$

Therefore any applied High objective makes the information system High overall. A genuinely Moderate system cannot contain a High system objective. But the overall word *High* does not replace the vector. If the category is H/M/L, PAIN needs H/M/L, not H/H/H.

DERIVATION

Use the high-water mark twice, for two different outputs.

1. Take a maximum per objective to retain the system vector.
2. Take the maximum of that vector to name the overall FIPS 199 impact level.

The overall label is useful for authorization and reporting. The dimensional vector is the input needed for an objective-sensitive PAIN calculation.

6 Certification Class as an operational assurance commitment

6.1 Class is not the consequence vector

FedRAMP's 2026 terminology uses Certification Classes A–D for package specifications. Current FedRAMP guidance describes those classes as levels of assurance that can support agency use cases; it does not make the Class letter a substitute for an agency's FIPS 199 C/I/A category. In particular, the guidance states that some Class C packages may provide adequate information for some High security objectives [FEDRAMP-CLASS].

This method therefore records Certification Class separately. Class can corroborate an intended-use assessment and selects FedRAMP processes outside the PAIN impact equation. Where confirmed intended-use evidence is incomplete, it can also inform a disclosed, rebuttable planning approximation of the ceiling. It does not by itself impose a categorical ceiling or flatten the vector. A Class C offering can therefore participate in an agency system whose applied vector contains a High objective. Under FIPS 199 that agency information system is High overall even though the offering's certification remains Class C.

That combination does not imply Class D-equivalent assurance. The agency consequence and provider commitment answer different questions:

Consequence vector
: What adverse effect could loss of C, I, or A have on the agency's mission, assets, or affected individuals?

Certification Class
: What reusable evidence, verification, monitoring, vulnerability-response, and incident-communication commitments does the provider supply under its certification?

6.2 A planning approximation, not a categorization rule

The former Low/Moderate/High shorthand remains useful as an initial approximation of assurance fit. Where confirmed intended-use evidence is incomplete, the same shorthand may also supply a provisional, rebuttable planning envelope for the ceiling. It must not be used as a substitute for the FIPS 199 vector, as an authorization rule, or in preference to actual-use evidence:

Class	Planning shorthand	Interpretation
A	Entry/non-sensitive	Limited reusable federal assurance; FedRAMP still describes some objective-specific uses above that shorthand.
B	Low	A reasonable default fit for most Low objectives, with agency-specific exceptions possible.
C	Moderate	A reasonable default fit for most Low and Moderate objectives; some High-objective uses remain possible.
D	High/broad	The broadest unclassified assurance fit and the tightest standard operational timeframes.

Table 3: A planning approximation for assurance fit, not a mapping from Class to a FIPS 199 security category.

The approximation is useful precisely because the Classes are progressive. It tells an agency where a service is likely to fit without claiming that every use at that impact level is supported or that every exception is prohibited.

6.3 Agency disclosure and provider acceptance

For any proposed Class B or Class C use that would place a High value in the intended-use ceiling, the deploying agency MUST affirmatively identify:

- the applicable information type or types;
- the objective or objectives rated High;
- the affected use and asset scope; and
- the SP 800-60 special factor, agency adjustment, or other governing basis.

The CSP MUST NOT infer that High use from an agency's identity, a generic text field, file upload, integration capability, hypothetical content, or an unrelated agency system. Once the agency communicates the intended use, the CSP MUST state whether it supports that use inside the offering's authorized and contractual scope. Technical ability to accept the bytes is not acceptance of the use.

The CSP MAY respond by accepting the existing Class commitment, adding controls, redesigning the service, offering tighter contractual response terms, or pursuing a higher Certification Class. It may also decline the use. The CSP cannot change the agency's impact determination, and the agency cannot unilaterally claim that a lower-Class service supplies higher-Class commitments.

The agency-identified data changes the ceiling; it does not by itself change the reusable asset security-impact profile. The profile changes only when the component's actual role, trust, privilege, dependency, data handling, or consequence inside the CSO changes. A provider redesign made to support the new use may justify such a change, but the new data label alone does not.

6.4 What higher assurance buys

FedRAMP makes the assurance difference operational. Table 4 shows selected current Class-dependent timeframes. The vulnerability timeframe is selected from Certification Class, PAIN, likely exploitability, and internet reachability. Incident communication is likewise selected from Class and PAIN [FEDRAMP-VDR, FEDRAMP-IEC].

Illustrative PAIN-5 commitment	**Class B**	**Class C**	**Class D**
LEV + IRV vulnerability mitigation/remediation	4 days	2 days	12 hours
Initial incident report	6 hours	1 hour	15 minutes
Ongoing incident reports	1 business day	6 hours	3 hours

Table 4: Selected 2026 operational timeframes. FedRAMP may revise the governing tables; implementations must use the current rules.

The incident values above govern communication, not a universal containment or recovery deadline. A contract can require tighter notification, containment, recovery, or mitigation commitments than the Class minimum. Such terms increase the assurance supplied to that customer; they do not by themselves convert the underlying certification to a higher Class.

6.5 The authorization and contract decision

Before a Class B or Class C offering is used for an agency-identified High objective, the agency MUST be told which Class-based timeframes apply. The agency and CSP should record one of four dispositions:

Standard assurance fit
: The agency determines that the existing Class commitments are adequate for the intended use.

Contractually enhanced
: The CSP accepts additional controls, evidence, or response commitments for the agency use.

Accepted assurance divergence
: The agency knowingly accepts the difference between the High consequence and the existing Class commitments.

Unsupported intended use
: The CSP does not accept the proposed information or use inside the offering's supported scope.

For recurring High-objective demand, a higher Certification Class is ordinarily the clearer and more reusable path. For a bounded use, contract terms and agency-managed controls may supply the needed assurance. In every case, the decision is an authorization and contracting matter. It must not be implemented by lowering the agency's C/I/A vector or by changing the resulting PAIN solely to make the existing service commitment appear sufficient.

6.6 Availability-specific assurance case

Availability deserves an explicit treatment because an applied A:H objective has direct implications for continuity engineering and service commitments. A:H means that loss of availability could have a severe or catastrophic adverse effect on the agency. It does not, by itself, prescribe multi-region

deployment, a particular uptime percentage, or universal Recovery Time Objective (RTO) and Recovery Point Objective (RPO) values.

The agency and CSP SHOULD translate the High consequence into an availability-assurance case that addresses, as applicable:

- mission time sensitivity, maximum tolerable disruption, and critical operating windows;
- the uptime commitment, its measurement boundary, exclusions, and dependent services;
- RTO and RPO by service or data tier;
- redundancy, failure-domain isolation, backup and restoration, failover mode, and degraded-service behavior;
- external and provider-managed dependencies that can defeat the design;
- recovery and failover test evidence, operating staffing, escalation, and agency communication; and
- any agency-managed continuity measures or residual risk.

This is an evidence discussion, not a checklist that makes any single mechanism dispositive. Multi-region deployment can still share a vulnerable software path, and a high uptime percentage does not alone establish recovery capability. The strength of the assurance case comes from the traceable relationship between the agency's mission tolerance and the combined architecture, operations, testing, and contractual commitments.

When agency-authorized information types raise the intended-use ceiling to A:H for a Class C offering, the existing uptime SLA and disaster-recovery strategy may support a *Standard assurance fit* disposition for the bounded use. If they do not, the parties can select *Contractually enhanced* terms, add architectural or agency-managed continuity measures, accept a documented assurance divergence, decline the use, or pursue Class D. The record SHOULD identify the chosen disposition and the evidence supporting it.

An adequate availability-assurance case can support use of a Class C offering without necessarily moving the offering to Class D. It does not reclassify a system whose FIPS 199 vector contains A:H, automatically lower the intended-use ceiling, or change the asset security-impact profile. It shows why the Class C level of assurance, including any supplemental commitments, is acceptable for that High-availability consequence. This recommendation is specific to Availability; it is not a claim that uptime or disaster-recovery measures satisfy High Confidentiality or Integrity objectives.

7 Intended use, not hypothetical capability

7.1 Two sets

Let:

T_{CSO}
: Information types the offering is designed and authorized to process, including applicable service-generated records needed to operate the offering.

T_a
: Information types agency a intends and is authorized to process through the offering.

The applicable set for that agency and offering is:

$$T_{a,\text{CSO}} = T_{\text{CSO}} \cap T_a. \tag{4}$$

Each entry is more than an identifier. It carries the agency-confirmed applied profile and the special-factor rationale for this use.

7.2 Generic input mechanisms

A text box can technically accept a medical narrative, criminal-justice record, tax return, trade secret, grocery list, or poem. That technical fact does not place all corresponding information types inside the intended boundary. The same is true of file uploads, object storage, email ingestion, and general APIs.

The CSO documents what it is designed and authorized to do. The agency documents what it intends and is authorized to put there. Only the confirmed overlap enters Equation (4). Prohibited information remains prohibited. Accidental or unauthorized data spillage is managed as spillage, not normalized into the ordinary categorization merely because the user interface could carry it.

This rule prevents generic product capability from becoming an automatic High designation. It also prevents the opposite error: a documented free-text, attachment, or integration workflow that is intended to carry a sensitive information type cannot be ignored just because the product has a generic label.

7.3 System-generated and control information

The intersection is not limited to agency-uploaded content. A CSO can generate audit records, identity material, security telemetry, configurations, decisions, transactions, or derived aggregates that require categorization. NIST SP 800-60 also directs reviewers to consider critical system functionality and security policy enforcement information.

The assessment therefore records:

- agency-supplied intended information;
- service-generated information;
- derived or aggregated information;
- trusted decisions and actions performed by the CSO;
- documented system-level adjustments.

These facts can alter an applied profile. They do not erase the requirement to identify the objective and rationale being changed.

8 From applied information types to a ceiling

8.1 One agency

For agency a, objective $o \in \{C, I, A\}$, and confirmed applied profile $q_{a,t}$, define:

$$E_a(o) = \max_{t \in T_{a,\mathrm{CSO}}} q^o_{a,t}. \tag{5}$$

Apply any documented system-level adjustment after the information-type maximum. The result

$$E_a = (E_a(C), E_a(I), E_a(A))$$

is the intended-use ceiling for that agency's deployment.

If no intended information types have been confirmed, the method does not invent a low ceiling. It records an unresolved assessment and uses a conservative, explicitly provisional upper bound until the intended use is confirmed.

IN PLAIN TERMS

For one agency, list the federal information types actually intended for this service. Then look at Confidentiality, Integrity, and Availability separately and keep the highest confirmed value in each column. A result such as M/H/L stays M/H/L; the overall FIPS 199 label High does not turn it into H/H/H. If the agency has not confirmed its intended information, do not guess Low. Mark the assessment unresolved and use a clearly provisional, conservative ceiling.

8.2 Several agencies and affected scope

Let A_x be the definite agencies whose information or mission can be affected through asset x. The asset-scoped ceiling is:

$$E_x(o) = \max_{a \in A_x} E_a(o). \tag{6}$$

This is deliberately asset-scoped:

- A single-tenant component receives the ceiling for its one definite agency.
- A shared control plane receives the per-objective maximum for every definite agency it can affect.
- Adding an agency to a shared component can raise, but never lower, its ceiling.
- Adding an agency elsewhere does not change a component that cannot affect that agency.

Agency identity alone does not determine an objective. An agency may operate many systems with different information types. Only its intended use of the assessed CSO participates.

IN PLAIN TERMS

For a shared asset, first identify the agencies that the asset can actually affect. Compare only those agencies' confirmed C/I/A ceilings and keep the highest value in each dimension. Do not include an agency merely because it is a CSP customer or uses another part of the offering. Sharing can widen the affected scope, but the number or identity of agencies does not raise a requirement unless their intended uses introduce a higher government consequence.

8.3 The vector approximation

Some implementations already maintain only two aggregate vectors:

- an offering or System Security Objectives vector (SSO);
- an Agency Security Objectives vector (ASO).

They can approximate Equation (5) with:

$$\widehat{E}_a(o) = \min(\mathrm{SSO}(o), \mathrm{ASO}_a(o)). \tag{7}$$

For several agencies, take the per-objective maximum of their ASO vectors before the minimum. This approximation is simple and safe when the aggregate vectors faithfully describe the same intended-use set. It can be imprecise when the CSO and agency vectors become High for unrelated information types. Two aggregate High values prove dimensional compatibility, not type-level overlap.

The preferred audit record therefore retains confirmed information-type membership. Equation (7) remains a documented fallback, not a claim that the set intersection was evaluated.

IN PLAIN TERMS

The shortcut compares two summary vectors: what the offering is designed to support and what the agency expects to use. It is useful when the underlying information-type records are unavailable, but it cannot prove that the two sides became High for the same reason. The type-aware method asks both, "How high can each side go?" and, "Are they talking about the same information and use?" Prefer the type-aware record. If the shortcut is used, label it as an approximation and retain that fact in the audit trail.

9 The asset security-impact profile remains separate

9.1 A profile, not a required taxonomy

The profile describes the uncapped confidentiality, integrity, and availability requirements of an asset inside the CSO. It is an input contract, not a required classification taxonomy or metadata key. A CSP MAY derive it by assigning the three objectives directly, by mapping a role-based archetype, or through another governed and deterministic method.

Role-based archetypes are a useful example because the same software can serve different roles:

- an in-memory store used as a disposable cache;
- the same store holding durable sessions or authentication tokens;
- the same store brokering mission-driving work;
- the same engine acting as a system-of-record database.

Those roles justify different uncapped profiles even before the agency context is applied. A role-based catalog is therefore system-relative: a provider derives it from the CSO architecture, trust relationships, control functions, and consequences of component compromise. The archetype is optional; the independently dimensional profile is not.

A conforming assignment MUST permit CR, IR, and AR to differ and SHOULD retain the derivation method, source evidence, objective-specific rationale, policy version, and any override. A scalar asset-value label may be used as a low-resolution shorthand for L/L/L, M/M/M, or H/H/H, but it SHOULD NOT be the only assignment interface because it forces every objective to move together and preserves less of the derivation evidence.

9.2 Effective requirements

Let the uncapped asset security-impact profile for asset x be

$$R_x = (R_x(C), R_x(I), R_x(A)).$$

The effective requirements handed to PAIN are:

$$R_x^{\text{eff}}(o) = \min(R_x(o), E_x(o)), \qquad o \in \{C, I, A\}. \tag{8}$$

The ceiling never raises a profile requirement. A High ceiling means only that High remains available where the asset's system-relative profile supports it. Conversely, a High profile requirement does not manufacture a High federal consequence when the intended-use ceiling is Moderate.

Input	Example	Scope	Meaning
Asset security-impact profile	H/H/M	Asset in the CSO	Uncapped system-relative consequence profile
Intended-use ceiling	M/M/L	Agency or affected agency set	Maximum federal consequence supported by confirmed intended use
Effective requirements	M/M/L	Asset in that agency scope	Per-objective minimum supplied to PAIN

Table 5: The three vectors remain visible in the audit chain.

9.3 Why this improves reuse

The same CSO can retain one governed profile derivation policy across several agency deployments. If it uses archetypes, it can reuse the catalog; if it uses direct per-objective assessment, it can reuse that evidence model. The customer-specific information appears in the ceiling. An architectural change still changes the profile. An intended-use change still changes the ceiling. Reviewers can challenge either judgment without rebuilding both.

Cross-system trust anchors, software-delivery paths onto agency devices, and shared CSP infrastructure require particular care because their consequences may extend beyond the information stored in the assessed component. The correct response is to model their architectural role and affected scope explicitly, not to assign a data-only profile and then hide an exception.

10 Worked derivations

10.1 Moderate intended use on a high-role database

A CSO's system-of-record database has an uncapped security-impact profile H/H/M, derived here from a role-based archetype. One agency's confirmed intended information types aggregate to M/M/L. Therefore:

$$R_{\text{db}}^{\text{eff}} = \min((H, H, M), (M, M, L)) = (M, M, L).$$

The database retains the same uncapped profile and role assignment. No customer-specific variant is created. If a later agency uses the same CSO differently, its ceiling changes without rewriting the database assessment.

10.2 High confidentiality without H/H/H flattening

An agency's applied information types and special factors aggregate to H/M/L. The overall FIPS 199 impact level is High because one objective is High. The ceiling remains H/M/L.

For a secrets service with an H/H/H security-impact profile:

$$R_{\text{secrets}}^{\text{eff}} = \min((H, H, H), (H, M, L)) = (H, M, L).$$

Turning the overall High label into H/H/H would discard the agency's dimensional judgment and overstate Integrity and Availability.

10.3 A shared asset across two agencies

Agency 1 has ceiling M/M/L. Agency 2 has ceiling L/H/M. A component dedicated to Agency 1 retains M/M/L. A component dedicated to Agency 2 retains L/H/M. A shared control plane that can affect both receives:

$$E_{\text{shared}} = \max((M, M, L), (L, H, M)) = (M, H, M).$$

Multi-agency scope increases only the objectives supported by a definite affected agency. It does not automatically force H/H/H.

10.4 A generic text field

A project-management CSO provides comments and attachments. Its documented intended use covers project plans, task records, and approved supporting artifacts. The agency has prohibited health, tax, criminal-justice, and classified content.

The text field does not add every NIST information type to T_{CSO}. Confirmed project and planning types enter the ceiling. Prohibited content remains governed by data-use policy and spillage response. If the agency later authorizes a new intended record type, the type inventory and ceiling are reviewed as a change.

10.5 Government impact and CSP reputation

An information type has applied C:L because disclosure would have limited adverse effect on the agency's operations. The CSP believes any public incident involving the record would materially harm its brand. The ceiling retains C:L unless the CSP can show how that reputational event creates a greater adverse effect on the agency, federal assets, or affected individuals. The CSP can and should manage its separate corporate exposure through enterprise risk controls.

11 Cross-CSP composition and the hyperscaler boundary

11.1 The knowledge problem

An upstream CSP may support many downstream CSOs without knowing which agencies ultimately depend on each service or which information types those agencies process. Requiring a hyperscaler to collect every downstream agency customer list is unlikely to be practical and may conflict with contractual boundaries.

The missing information does not prove a low ceiling. It creates uncertainty about the dimensional aggregate.

11.2 Privacy-preserving vector attestations

A downstream CSO could share a bounded assertion without naming its agency customers or detailed information types:

- downstream CSO identifier;
- authorization or categorization boundary;
- maximum required C/I/A ceiling for the dependency;
- whether the dependency is single-customer or shared;
- evidence status, effective date, and expiry;

- attesting party and signature.

The upstream CSP can take the per-objective maximum of current assertions. This is less precise than type-level evidence but more informative than customer identities alone.

11.3 A non-normative saturation approximation

Where no downstream vector assertions exist and the upstream service supports a broad, heterogeneous population of confirmed downstream CSOs, the dimensional maximum can reasonably be expected to tend toward the upper envelope used for the upstream system category.

Let x be the number of confirmed downstream CSPs whose CSOs depend on the upstream service and whose agency-specific intended-use vectors are unavailable to the upstream CSP. Let $\widehat{E}_{\text{downstream}}(x)$ denote the conservative policy approximation, not an observed population statistic. The asymptotic claim is:

$$\lim_{x\to\infty} \widehat{E}_{\text{downstream}}(x) = \begin{cases} (M, M, M), & \text{Moderate-impact upstream system,} \\ (H, H, H), & \text{High-impact upstream system.} \end{cases}$$

This limit states only the conservative endpoint of the approximation. It does not claim that downstream uses are independent samples, specify a distribution over information types, measure a probability of convergence, or define a function for finite x. This paper makes no claim, approximation, or recommendation about a dependency count, diversity measure, confidence level, or other point at which an upstream CSP could treat a finite population as having reached the endpoint.

The observation is narrower: for an upstream CSP supporting an extremely broad, heterogeneous population of downstream CSPs while lacking their agency-specific vectors, the applicable upper-envelope vector is a reasonable conservative approximation of the otherwise unknowable aggregate. Whether to use that approximation for any finite deployment is an adopter-specific risk and governance decision outside this paper. Any use must be disclosed as an assumption rather than presented as a measured result.

Capability and justified finite decisions. Nothing in the limit prevents an upstream CSP from making a supportable dimension-by-dimension decision for its actual finite population. Where the CSP can document the expected and permitted downstream information types, prohibited uses, service boundary, customer commitments, and relevant technical evidence, it may justify setting one or more dimensions below the asymptotic envelope. That decision is derived from the documented use case, not from a claim that a particular value of x proves convergence. Its assumptions and reassessment conditions should remain in the governance record.

Large hyperscalers also illustrate why technical capability and government consequence must remain separate. AWS documents FedRAMP High-authorized services in AWS GovCloud (US), and Google documents a FedRAMP High-authorized cloud boundary and in-scope services [AWS-HIGH, GCP-HIGH]. Those defined scopes demonstrate that some upstream providers already possess controls and operational capabilities that can support some High-objective uses. They do not make every hyperscaler service H/H/H, authorize every downstream information type, or determine the FIPS 199 category of an agency system. Capability is evidence about the assurance a service can supply; the intended-use vector still states the government consequence at stake.

The approximation never authorizes a use outside the upstream boundary. If a downstream use actually requires High while the upstream information system is genuinely Moderate, the result is a boundary mismatch, not permission to process the use.

SCOPE BOUNDARY

This paper identifies the cross-CSP composition problem but does not standardize a saturation threshold, convergence rule, or attestation protocol. It offers an asymptotic upper-envelope observation, not guidance for deciding when a finite deployment may use it. Developing a complete supply-chain profile is an exercise left to the reader.

12 Governance and audit record

A deterministic equation does not make uncertain inputs certain. A conforming assessment preserves evidence and ownership for each judgment.

12.1 Minimum record

The record SHOULD contain:

- CSO identity, boundary, designed functions, and service-generated data;
- confirmed intended and prohibited information types;
- per-agency information types, relationship (definite or target), and governing evidence;
- each provisional and applied C/I/A profile;
- special factors considered and objective-specific adjustment rationale;
- system-level aggregation or critical-function adjustments;
- the per-agency and asset-scoped ceilings;
- the overall FIPS 199 impact label, stored separately from the vector;
- the uncapped asset security-impact profile, its derivation method and rationale, the effective requirements, and every capped objective;
- multi-agency affected-scope basis;
- confidence, assumptions, manual-review items, effective date, and expiry;
- any vector approximation or downstream-CSP saturation approximation used.

Operator attestations remain separate from analyst or automated inferences. Confidence measures evidence quality and never lowers an objective. When two adjacent values remain credible, select the higher value provisionally and state what evidence would justify lowering it.

12.2 Change triggers

Reassess the ceiling when:

- an agency adds or removes an intended information type;
- a special factor becomes applicable or ceases to apply;
- a CSO adds a new ingestion path, generated record, trusted decision, or system function;
- tenancy or shared-control-plane scope changes;

- aggregation creates a new system-level consequence;
- the governing categorization or authorization boundary changes;
- a downstream-CSP assertion expires or the known dependency composition changes materially.

Security-impact-profile changes follow different triggers: component role, trust, privilege, dependency, data handling, or architectural consequence. Keeping the triggers separate is one of the model's principal governance benefits.

13 Handoff contract

The output of this paper is not a PAIN word or N-rating. It is an auditable, three-dimensional input:

DERIVATION

For each affected asset x:

$$\boxed{R_x^{\text{eff}}(o) = \min\left(R_x(o), \max_{a \in A_x} \max_{t \in T_{\text{CSO}} \cap T_a} q_{a,t}^{o}\right)}$$

for $o \in \{C, I, A\}$, with documented system-level adjustments applied where required.

A transport may encode the resulting vector as:

```
CR:M/IR:H/AR:L
```

or:

```
cr-m_ir-h_ar-l
```

The handoff must also identify the uncapped asset security-impact profile, its derivation method and source, the ceiling vector, ceiling source and scope, and whether any objective was capped. The downstream PAIN method then combines this effective requirements vector with vulnerability impact.

If a Class B or Class C intended use contains a High objective, the adjacent governance record must also identify the agency attestation, Certification Class, assurance disposition, and any contract or agency-managed safeguards. These fields do not modify R_x^{eff}; they explain whether the operational commitment is sufficient for the consequence the vector represents.

14 How the method may—and may not—be used

This method is intentionally narrow. It answers one question: given the government information use that has actually been identified and the role an asset actually performs, what is the greatest C/I/A requirement that PAIN should receive for that asset? The answer is an effective Security Requirements vector with a traceable derivation. It is not a second authorization decision.

14.1 The output boundary

The method produces R_x^{eff}, a per-asset CR/IR/AR input to the PAIN equation. That vector is not, by itself:

- the agency's FIPS 199 system category or control baseline;

- permission for the CSO to accept a particular information type;
- the CSO's FedRAMP Certification Class;
- a statement that the available assurance is adequate for the intended use; or
- a PAIN word, N-level, remediation deadline, or incident-response deadline.

Those decisions consume different evidence and occur at different points in the authorization and vulnerability-response processes. Keeping them separate prevents a convenient operational label from silently changing the government consequence being measured.

14.2 The authority and evidence boundary

NIST SP 800-60 provides provisional, government-centered impact recommendations. The deploying agency remains responsible for identifying the information types authorized for the system, applying special factors and governing sources, and confirming the resulting FIPS 199 categorization. The catalog does not make an agency determination merely because a type appears in it, and the counts in this paper describe the catalog's text rather than the prevalence of High-impact agency deployments.

The intersection of agency use and supported CSO use narrows the relevant information types; it does not eliminate review of system-level aggregation, critical functions, control information, or extenuating circumstances. Classified information and national-security systems remain outside this paper's scope.

When only a vector approximation of an information-type set is available, the approximation can combine objective maxima that came from different underlying types. It is suitable as a conservative envelope, but it cannot be cited as proof that one particular information type has that combined profile. The approximation must be labeled and its source evidence retained.

14.3 The ceiling boundary

The intended-use vector is a cap, not a floor and not a product-capability score. Applying it can reduce a profile requirement that exceeds the consequence supported by the agency's intended use; it cannot raise a low-requirement asset merely because the system contains more sensitive information. The uncapped security-impact profile continues to describe the reusable asset assessment, while the effective vector describes that asset in this system and agency context.

A lower effective requirement on one asset does not lower the FIPS 199 category of the system, remove applicable controls, or authorize a broader data use. It only bounds the consequence input used to evaluate a vulnerability affecting that asset. Likewise, a generic text field, upload mechanism, or technical ability to receive bytes does not add hypothetical information types to the ceiling. Accidental or prohibited data spillage remains a security and governance event; it is not a reason to categorize every generic input for every imaginable use.

The inverse is equally important. Once an agency identifies and justifies an authorized High objective, the CSP must not lower that objective merely to make it look consistent with a Certification Class shorthand. If the offering cannot support the use at an acceptable level of assurance, the result is an authorization, architecture, contract, or acceptance problem—not a reason to rewrite the consequence.

14.4 The Certification Class and assurance boundary

Certification Class and the intended-use vector answer different questions. Class describes the reusable assurance and operational commitments supplied by the CSP; the vector describes the adverse effect the government could experience. Class is therefore not a data-eligibility label, FIPS 199

category, or impact ceiling by itself. It may inform a provisional ceiling approximation when actual-use evidence is incomplete; the Low/Moderate/High Class shorthand remains a rebuttable planning aid rather than a categorization rule.

For a Class B or Class C use containing a High objective, the agency must identify and justify that objective and the CSP must affirmatively accept the use. Neither agency identity nor generic technical capability supplies that agreement. A Class C offering can support an agency use with a High objective, but it continues to supply Class C commitments unless a contract or higher Class provides more. The agency and CSP must record whether that level of assurance is adequate rather than altering the impact vector to conceal a divergence.

FedRAMP's published timeframes can change, and contracts can supplement them. Incident-reporting timeframes are communication obligations, not universal containment or recovery SLAs. Availability also requires a specific distinction: A:H states a severe or catastrophic government consequence, but it does not itself prescribe an uptime percentage, multi-region topology, RTO, or RPO. Architecture, testing, recovery evidence, agency continuity measures, and contract terms can establish an acceptable availability-assurance case; they do not recategorize the system or lower its effective AR.

14.5 The downstream PAIN boundary

This paper ends when it hands R_x^{eff} to PAIN. Under the reference calibration in [PAIN], an effective vector containing only Low and Moderate requirements cannot produce a Debilitating result. Combined with the FIPS 199 high-water-mark rule, this means a genuinely Low- or Moderate-impact system cannot produce Debilitating under that calibration. This is a consequence of the PAIN formula and thresholds; NIST SP 800-60 and FIPS 199 do not independently classify vulnerabilities as Debilitating.

IN PLAIN TERMS

In plain terms, *Debilitating* is PAIN's catastrophic customer-effect class. FedRAMP introduced the term as the replacement for *Catastrophic Adverse Effect* [FEDRAMP-PAIN]. FIPS 199 and NIST SP 800-60, in turn, assign High when loss of a security objective could have a severe or catastrophic adverse effect on organizational operations, organizational assets, or individuals [FIPS199, SP800-60-1]. Requiring an effective High objective before Debilitating becomes reachable therefore aligns PAIN's highest single-agency consequence class with NIST's High-impact consequence category.

A High objective changes reachability, not the result for every vulnerability. The reference calibration still requires at least one High technical impact aligned with a High effective requirement, plus an additional High technical impact aligned with a Moderate-or-High effective requirement. A single High/High alignment remains Disruptive under that calibration. Consequently, a Debilitating result in a Class B or Class C offering is exceptional but not impossible: the agency-identified High objective must survive the asset cap and the vulnerability must satisfy the compound alignment.

Certification Class does not mechanically determine that PAIN word, but it can affect PAIN indirectly. Where exact intended-use evidence is incomplete, the Class planning shorthand can inform a reasonable, disclosed, and rebuttable approximation of the ceiling for the information and use the offering is expected to support. Because that ceiling caps R_x^{eff}, the approximation constrains which PAIN results are reachable.

The approximation is not conclusive. Confirmed agency information types, special factors, and accepted intended use displace it: an agency-identified High objective can remain High in a supported Class C use, while Class D does not raise a Low-role asset or turn every objective High. Once the consequence and affected scope are derived, Class also selects the standard level of assurance and operational timeframes. Providers may govern different PAIN calibrations, but any different thresholds or predicates must be disclosed as policy choices rather than attributed to FIPS 199, NIST SP 800-60, or Certification Class.

14.6 The cross-CSP approximation boundary

The hyperscaler saturation limit is a conservative upper-envelope policy, not a measured convergence result, probability model, or federal standard. It defines no finite dependency count, customer-diversity measure, or confidence threshold. A provider may make a justified dimension-by-dimension decision from evidence about its actual bounded population, but it cannot claim that a particular value of x proves the asymptotic endpoint has been reached.

15 Conclusion

The first input to an agency-specific PAIN calculation should be the federal consequence that is actually possible in the assessed use. NIST SP 800-60 supplies provisional, objective-specific evidence; FIPS 199 supplies dimensional aggregation and the overall system impact rule. Neither calls for replacing the resulting vector with one repeated High, Moderate, or Low label.

Separating the intended-use ceiling from the asset security-impact profile makes both judgments clearer. The profile records the component's system-relative requirements inside the CSO, whether derived directly, through an archetype, or by another governed method. The ceiling records what the definite agency use puts at stake. Their per-objective minimum becomes the effective CR/IR/AR vector handed to PAIN. Multi-agency components aggregate only the agencies they can affect. Generic technical capability does not authorize hypothetical data. CSP corporate consequences remain important but do not silently redefine government impact.

The result is a reusable, independently dimensional asset profile joined to an agency-specific, source-traceable federal consequence vector. Certification Class remains beside that vector as the provider's operational assurance commitment. When a High objective is proposed for a Class B or Class C offering, agency disclosure, CSP acceptance, the applicable response clocks, and any supplemental contract terms make the divergence explicit. The PAIN equation can then begin with its most consequential assumption made explicit without pretending that impact and assurance are the same thing. For A:H specifically, an availability-assurance case connects that consequence to uptime, recovery, architecture, testing, dependencies, and contract evidence without treating any single implementation as mandatory.

References

[FIPS199] National Institute of Standards and Technology, *Standards for Security Categorization of Federal Information and Information Systems*, FIPS PUB 199, February 2004. https://doi.org/10.6028/NIST.FIPS.199.

[SP800-60-1] K. Stine, R. Kissel, W. Barker, J. Fahlsing, and J. Gulick, *Guide for Mapping Types of Information and Information Systems to Security Categories, Volume I*, NIST SP 800-60 Volume I Revision 1, August 2008. https://doi.org/10.6028/NIST.SP.800-60v1r1.

[SP800-60-2] K. Stine, R. Kissel, W. Barker, A. Lee, J. Fahlsing, and J. Gulick, *Guide for Mapping Types of Information and Information Systems to Security Categories, Volume II*, NIST SP 800-60 Volume II Revision 1, August 2008. https://doi.org/10.6028/NIST.SP.800-60v2r1.

[CVSS31] FIRST, *Common Vulnerability Scoring System v3.1: Specification Document*, Environmental Security Requirements. https://www.first.org/cvss/v3.1/specification-document.

[FEDRAMP-VER] FedRAMP, *Vulnerability Evaluation and Reporting*, Consolidated Rules for 2026, rule VER-EVA-EPA, launched June 24, 2026. https://www.fedramp.gov/2026/providers/20x/rules/vulnerability-evaluation-and-reporting/.

[FEDRAMP-PAIN] FedRAMP, *Initial Outcome of RFC-0031: Updated Incident Communications Procedures*, Notice 0012, PAIN and customer-effect terminology. https://www.fedramp.gov/notices/0012/.

[FEDRAMP-CLASS] FedRAMP, *Choosing a Certification Class*, Consolidated Rules for 2026. https://www.fedramp.gov/2026/providers/start/class/#alignment-with-agency-use-cases.

[FEDRAMP-2026] FedRAMP, *What's Changing in 2026*, Consolidated Rules for 2026. https://www.fedramp.gov/2026/providers/updating/changes/.

[FEDRAMP-VDR] FedRAMP, *Vulnerability Detection and Response*, Consolidated Rules for 2026, rule VDR-TFR-PVR. https://www.fedramp.gov/2026/providers/20x/rules/vulnerability-detection-and-response/.

[FEDRAMP-IEC] FedRAMP, *Incident Evaluation and Communication*, Consolidated Rules for 2026, rules IEC-CSO-IIR and IEC-CSO-OIR. https://www.fedramp.gov/2026/providers/20x/rules/incident-evaluation-and-communication/.

[AWS-HIGH] Amazon Web Services, *FedRAMP Compliance*, AWS GovCloud (US) FedRAMP High scope and services. https://aws.amazon.com/compliance/fedramp/.

[GCP-HIGH] Google Cloud, *FedRAMP Compliance*, FedRAMP High cloud boundary and in-scope services. https://cloud.google.com/security/compliance/fedramp.

[PAIN] M. Venne, *A Deterministic, CVSS–Environmental Method for FedRAMP VDR/VER Vulnerability Prioritization*, stackArmor, July 2026. https://stackarmor.github.io/rfc-fedramp-vdr/vdr-pain-cvss.pdf.

Reproducibility record

The catalog analysis uses the source-pinned NIST SP 800-60 Volume II Revision 1 conversion distributed with `trivy-plugin-vdr-skills`. A conforming run must return:

Check	Expected result
Total catalog records	170
Scored information types	168
Non-scoring delivery mechanisms	2
Provisional profile contains at least one High	11
Provisional H/H/H	4
Additional records with explicit conditional High language	94
Unique provisional or conditional High path	105
No explicit provisional or conditional High path	63

The conditional count includes a record only when objective-specific `specialFactors` text expressly permits a High impact level. A conditional record contributes no High objective to an assessed system unless the agency confirms that the stated factor applies.