



C 4 S A M

C4SAM · WHITEPAPER

# TRUE RISK PRIORITY

From vulnerability findings to accountable remediation

How C4SAM combines threat evidence, exposure, technical impact, business context, and change execution in a model aligned with CISA BOD 26-04 decision principles.

| AUDIENCE                   | SCOPE                       | VERSION         |
|----------------------------|-----------------------------|-----------------|
| CIO · CISO · IT Operations | Executive & operating model | 1.0   July 2026 |

## POSITIONING STATEMENT

C4SAM True Risk Priority is a decision and execution layer. It does not replace scanners, CVSS, CISA KEV, CMDB, BIA, or IT service management. It turns those sources into a transparent remediation order and carries the decision into controlled change.

C4SAM GmbH

# Contents

---

|                                                              |           |
|--------------------------------------------------------------|-----------|
| <b>Document purpose and important qualification.....</b>     | <b>3</b>  |
| <b>1. Executive summary.....</b>                             | <b>4</b>  |
| 1.1 Benefits by stakeholder .....                            | 4         |
| <b>2. Why conventional prioritization breaks down .....</b>  | <b>6</b>  |
| 2.1 Failure modes of scanner-only queues .....               | 6         |
| 2.2 Why prioritization must be explainable .....             | 6         |
| <b>3. C4SAM True Risk Priority .....</b>                     | <b>8</b>  |
| 3.1 Weighted score .....                                     | 8         |
| 3.2 Escalation before thresholds .....                       | 8         |
| 3.3 Default priority bands.....                              | 9         |
| <b>4. Exposure and business context.....</b>                 | <b>10</b> |
| 4.1 Public zone is not the same as internet exposed .....    | 10        |
| 4.2 CMDB, BIA, and service relationships.....                | 11        |
| 4.3 Fallback rules and data quality .....                    | 12        |
| <b>5. CISA-inspired decision path.....</b>                   | <b>13</b> |
| 5.1 Factor completeness is a control .....                   | 13        |
| 5.2 BOD-style timeline examples .....                        | 14        |
| <b>6. From vulnerability to remediation .....</b>            | <b>16</b> |
| 6.1 End-to-end workflow.....                                 | 16        |
| 6.2 Instance-level integrity .....                           | 16        |
| 6.3 Change creation and CI matching.....                     | 16        |
| 6.4 Production integration principles.....                   | 18        |
| <b>7. The operating model.....</b>                           | <b>19</b> |
| 7.1 Roles and responsibilities.....                          | 19        |
| 7.2 Decision cadence.....                                    | 19        |
| 7.3 Exception and risk-acceptance control.....               | 20        |
| <b>8. Governance and control framework .....</b>             | <b>21</b> |
| 8.1 Overrides .....                                          | 21        |
| 8.2 Data freshness .....                                     | 21        |
| 8.3 Model limitations .....                                  | 21        |
| <b>9. Implementation architecture .....</b>                  | <b>23</b> |
| 9.1 Minimum viable data contract .....                       | 23        |
| 9.2 Event-driven recalculation .....                         | 23        |
| 9.3 Reliability and security requirements .....              | 23        |
| <b>10. Implementation roadmap.....</b>                       | <b>24</b> |
| 10.1 Pilot design.....                                       | 24        |
| 10.2 Adoption risks and mitigations .....                    | 24        |
| <b>11. Measurement and executive reporting .....</b>         | <b>26</b> |
| 11.1 Executive dashboard narrative .....                     | 26        |
| <b>12. Illustrative decision scenarios.....</b>              | <b>28</b> |
| <b>13. Recommendations for CIOs and CISOs.....</b>           | <b>29</b> |
| <b>A. Appendix: field and evidence model .....</b>           | <b>30</b> |
| <b>B. Appendix: policy checklist .....</b>                   | <b>31</b> |
| <b>C. References and source notes.....</b>                   | <b>32</b> |
| <b>D. Appendix: glossary of terms and abbreviations.....</b> | <b>33</b> |
| <b>About C4SAM.....</b>                                      | <b>35</b> |

## Document purpose and important qualification

---

This whitepaper explains the C4SAM True Risk Priority approach for senior executives and operational leaders. It describes the business problem, the decision model, the data and governance requirements, and the integration of vulnerability management with remediation through change management.

**Important terminology** On June 10, 2026, CISA published Binding Operational Directive (BOD) 26-04, “Prioritizing Security Updates Based on Risk.” It supersedes BOD 19-02 and BOD 22-01 and prioritizes remediation using four decision criteria: public exposure, KEV status, exploit automation, and technical impact. The C4SAM True Risk method is aligned with these criteria. Note the legal scope: BOD 26-04 is binding only for U.S. Federal Civilian Executive Branch agencies. For all other organizations, C4SAM timelines and SLAs remain configurable customer policy and must not be represented as an official CISA compliance determination.

The model and default thresholds described here represent the July 2026 C4SAM implementation and can be governed as customer policy. Production adoption requires validated asset, exposure, BIA, service, ownership, and vulnerability-enrichment data.

# 1. Executive summary

The decision that vulnerability programs actually need

Most organizations do not lack vulnerability findings. They struggle to turn a large, technically ranked backlog into a clear and defensible sequence of business actions. Scanner severity and CVSS remain essential descriptions of technical weakness, but they do not reliably answer four executive questions: Which finding threatens us most? On which exact asset or service? Why must it move ahead of other work? Is remediation being executed within an accountable change process?

C4SAM True Risk Priority addresses that gap by combining four evidence domains into one explainable decision: threat, exposure, technical impact, and business impact. It then applies explicit escalation rules before score thresholds, producing a priority from P1 to P5, a reason trail, and a policy-based remediation target. The result is more than another score. It is a shared working agreement between security, infrastructure, application owners, business service owners, and change management.

| Executive concern                   | C4SAM response                                                                       | Expected outcome                                      |
|-------------------------------------|--------------------------------------------------------------------------------------|-------------------------------------------------------|
| Too many critical findings          | Contextual ranking by exploit evidence, reachability, impact, and service importance | A smaller, defensible urgent queue                    |
| Security and operations disagree    | Visible reason codes, factor status, thresholds, and escalation rules                | A shared basis for remediation decisions              |
| Findings do not become work         | Change-request creation linked to the affected Cockpit object and iTop CI            | Traceable transition from risk to execution           |
| Executive reporting counts activity | Priority, due date, ownership, change status, and aging                              | Evidence of risk reduction rather than ticket volume  |
| Context is incomplete               | Explicit fallback path and missing-data indicators                                   | Useful decisions without false claims of completeness |

Core message for leadership True Risk Priority changes the management question from “How many critical vulnerabilities do we have?” to “Which exposures can materially harm the organization, who owns the response, what change is underway, and when will residual risk be reduced?”

## 1.1 Benefits by stakeholder

| Stakeholder                    | Primary benefit                                                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------|
| CIO                            | A portfolio view that aligns cyber remediation with service continuity, change capacity, and technology lifecycle. |
| CISO                           | A transparent risk-prioritization method grounded in exploit evidence, exposure, and organizational impact.        |
| IT operations                  | A smaller queue with clear rationale, affected configuration items, remediation text, urgency, and due dates.      |
| Service and application owners | Visibility into how service criticality, BIA, data classification, and ownership affect priority.                  |
| Risk and audit                 | A reproducible decision trail that separates complete CISA-style inputs from customer-policy fallback scoring.     |

## 2. Why conventional prioritization breaks down

---

Technical severity is necessary, but operationally incomplete

A vulnerability score describes properties of a weakness under defined assumptions. An enterprise remediation decision must additionally consider whether the affected instance is reachable, whether exploitation is occurring in the wild, what control an attacker could gain, what service depends on the asset, how much harm interruption would cause, and whether an accountable owner can execute the fix.

This difference matters in practice. The same vulnerability can be P1 on an internet-reachable production endpoint, P4 on a restricted internal system, and P5 on an isolated lab instance. In contrast, a moderate scanner finding can require urgent action when it is publicly reachable, automatable, or linked to a high-value business service. A single global severity label cannot represent these differences.

### 2.1 Failure modes of scanner-only queues

- Volume inflation: many findings are labelled high or critical, overwhelming limited remediation capacity.
- Asset blindness: a weakness is ranked without understanding the network zone, publication path, or service relationship of the affected instance.
- Threat blindness: theoretical exploitability is confused with evidence of exploitation in the wild.
- Business blindness: a development host and a revenue-critical production service can receive the same apparent urgency.
- Ownership ambiguity: findings with no clear owner remain in reports but do not become controlled work.
- Workflow discontinuity: vulnerability tools create alerts while IT operations works in a separate change system, breaking traceability.
- Metric distortion: teams optimize closure counts or average age while the most damaging exposures remain open.

### 2.2 Why prioritization must be explainable

A risk model influences production change, downtime, emergency approvals, and resource allocation. Therefore, each result must be explainable in plain language. C4SAM records reason codes and component scores so that users can see whether priority was driven by KEV status, internet exposure, public-zone classification, automatable exploitation, total system control, business criticality, regulatory scope, or an ownership gap.

Explainability also protects against model drift. If the organization changes a threshold or weight, the effect can be previewed against fixed scenarios before activation. This supports policy governance without hiding decisions inside an opaque algorithm.

Back Home / Vulnerability / True Risk Rules Audit Mode Service 98%

Advanced True Risk scoring configuration. Falback rule lifecycle management is handled in the Rule Library.

**Mode**

- Legacy mode
- True Risk PoC mode**
- Manual overrides replace calculated True Risk

**Scoring Weights**

|                  |                 |
|------------------|-----------------|
| Threat           | Exposure        |
| 0.4              | 0.25            |
| Technical Impact | Business Impact |
| 0.15             | 0.2             |
| <b>Total LRP</b> |                 |

**Priority Rules**

|          |                  |
|----------|------------------|
| P2 score | P3 score         |
| 85       | 65               |
| P4 score | Exploit P2 score |
| 40       | 75               |

- KEV + Internet + Total Control => P1
- KEV + Internet + Automatable => P1
- Internet + Public Exploit + RCE => P1

**Priority SLAs**

P1 SLA  
Immediate action / 3 days / forensic triage required if exploited or total-control risk

P2 SLA  
Accelerated remediation / 14 days

P3 SLA  
Planned remediation / 30 days

P4 SLA  
Standard maintenance window / 60 days

P5 SLA  
Need regular upgrade cycle or accepted risk review

**CISA BOD 26-04 Mapping**

Full BOD 26-04 calculation uses four official decision points: Publicly Exposed, In KEV, Automatable, and Technical Impact.

**Full CISA BOD 26-04 calculation**  
CVE exists and all four decision points are available from security-zone exposure, KEV, and CVESSVC enrichment.

**Alternative True Risk fallback**  
Used when there is no CVE or CISA enrichment is missing. The finding still receives a True Risk Priority and SLA, but it should not be labelled BOD-compliant.  
Public exposure is set on Security Zone local BIA fields. Automatable and technical impact can come from CISA Vulnrichment/CVE SSVC data when available.

**Rule Impact Preview**

Shows how the current rule configuration changes priority and score for fixed demo vulnerabilities. Use it as a quick sanity check after changing weights, thresholds, or escalation toggles.

|                                             |    |          |
|---------------------------------------------|----|----------|
| Internet-facing KEV with total control      | P1 | Score 85 |
| Critical ERP service, Internal only         | P2 | Score 71 |
| Isolated lab RCE without exploit            | P3 | Score 38 |
| DMZ web server with public exploit          | P1 | Score 84 |
| Management system with high business impact | P2 | Score 65 |

Figure 1 – C4SAM True Risk configuration: weights, thresholds, escalation rules, SLA text, CISA-style status, and impact preview.

### 3. C4SAM True Risk Priority

A contextual decision layer over existing evidence

C4SAM calculates True Risk at the affected-instance level whenever context is available. It enriches vulnerability evidence with the asset and service graph, security-zone data, environment and BIA metadata, business and application services, data classification, regulatory scope, and ownership information. This allows one vulnerability definition to receive different priorities across different operational contexts.

#### 3.1 Weighted score

Default formula True Risk Score = Threat × 40% + Exposure × 25% + Technical Impact × 15% + Business Impact × 20%

| Component        | Default weight | Representative inputs                                                                      | Decision meaning                                      |
|------------------|----------------|--------------------------------------------------------------------------------------------|-------------------------------------------------------|
| Threat           | 40%            | Scanner/C4SAM score, CISA KEV, exploit evidence, public exploit, ransomware indicators     | How credible and active is the threat?                |
| Exposure         | 25%            | Publicly exposed, internet facing, zone exposure level and score, DMZ/management signals   | Can an attacker reach or traverse to the instance?    |
| Technical impact | 15%            | RCE, automatable exploitation, partial or total control, SSVC/Vulnrichment evidence        | What control or consequence can exploitation produce? |
| Business impact  | 20%            | BIA score, criticality, RTO/RPO, service importance, data classification, regulatory scope | How much organizational harm could result?            |

Weights are customer-governed policy values and must sum to 100%. They make the model adjustable to organizational risk appetite, but adjustment should be controlled, tested against scenarios, approved, and versioned. Frequent tuning to make individual findings “look right” undermines comparability and auditability.

#### 3.2 Escalation before thresholds

Not every risk relationship is safely represented by a weighted average. C4SAM therefore evaluates high-confidence escalation conditions before applying numerical thresholds. In the July 2026 implementation, explicit public exposure combined with internet-facing reachability invokes P1. Other configured examples include KEV plus internet exposure plus total control, KEV plus internet exposure plus automatable exploitation, and internet exposure plus public exploit plus remote code execution.

Why escalation comes first A weighted average can weaken a decisive signal. A publicly reachable, actively exploited route to total system control should not fall below urgent action because another data field is low or missing.

#### 3.3 Default priority bands

| Priority | Decision intent         | Default trigger after escalation                            | Operational treatment                                                                     |
|----------|-------------------------|-------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| P1       | Immediate action        | Explicit escalation condition                               | Incident-aware remediation, emergency or expedited change, forensic triage when warranted |
| P2       | Accelerated remediation | Score $\geq 85$ ; exploit evidence may qualify at $\geq 75$ | Prioritized change within a short policy window                                           |
| P3       | Planned remediation     | Score $\geq 65$ or ownership gap floor                      | Named owner and scheduled change                                                          |
| P4       | Standard maintenance    | Score $\geq 40$                                             | Normal maintenance window                                                                 |
| P5       | Lifecycle or review     | Score $< 40$ without escalation                             | Next upgrade cycle, compensating control, or accepted-risk review                         |

Priority is a management decision, not a prediction of exact loss. It should direct attention, service-level policy, approval path, and reporting. The score provides ordering within and across priority bands; reason codes provide the justification.

## 4. Exposure and business context

Where a vulnerability exists changes what it means

### 4.1 Public zone is not the same as internet exposed

A public zone is an architecture classification. It describes where an asset belongs in the intended trust model, for example Public, DMZ, External, or Internet Edge. Internet exposed is evidence of reachability: an external actor can potentially connect through a public address, DNS name, load balancer, reverse proxy, VPN publication, or another externally reachable path.

| Zone classification | Reachability     | Interpretation                                                       |
|---------------------|------------------|----------------------------------------------------------------------|
| Public              | Internet exposed | Highest-confidence exposure signal; immediate escalation can apply.  |
| Public              | Not confirmed    | Elevated architectural context, but no confirmed external path.      |
| Internal            | Internet exposed | Reachability overrides stale or misleading topology classification.  |
| Internal            | Not exposed      | Internal baseline; lateral movement and service impact still matter. |

C4SAM prefers explicit reachability and source-record metadata. Configured fallback rules are used where direct fields are incomplete. This ordering is important: inferred labels should not override stronger evidence.

Edit

Name\*

public\_zone

Display Name\*

Public Zone

Criteria\*

Accessible to the public, includes public-facing services like the company website.

Controls Description\*

Basic firewall, intrusion detection system (IDS).

Must be 8-50 words long.

Local BIA / BCM Data ?

Security Zone Exposure Level

Internet

Security Zone Exposure Score

90

Internet Facing

Publicly Exposed (CISA BOD 26-04)

Data Classification

Confidential

Regulatory Scope

NIS2

Close

Save changes

Figure 2 – Security-zone context in C4SAM, including exposure level and score, internet-facing status, publicly exposed status, data classification, and regulatory scope.

4.2 CMDB, BIA, and service relationships

Business impact is not derived from vulnerability severity. It is derived from the role of the affected technology in the organization. C4SAM can use environment BIA criticality and impact score, RTO/RPO tiers, service and application criticality, data classification, regulatory scope, and the existence or absence of accountable ownership.

Update Environment

Name

Demo

Description

Demo Environmnet

Local BIA / BCM Data

BIA Business Criticality

Medium

BIA Business Impact Score

50

Data Classification

Internal

Regulatory Scope

NIS2

BIA RTO Tier

BIA RPO Tier

Close

Update

Figure 3 – Environment-level BIA and regulatory metadata provide a governed source for business-impact scoring.

Context should be inherited carefully. A security-zone rule can provide exposure defaults; an environment can provide continuity expectations; a business service can provide organizational criticality; and an application can provide data and publication context. C4SAM records the source of a decision so inherited defaults remain distinguishable from direct evidence.

4.3 Fallback rules and data quality

Fallback rules allow the program to operate while source systems mature. They can infer exposure or impact from recognized zone, environment, service, or application names. Fallback is useful, but it must be visible and governed: each rule needs a scope, owner, status, rationale, review date, and test scenario. Direct source metadata should take precedence when it becomes available.

True Risk Rule Library

Save active rule

Revert draft

Fallback Rule Inventory

Manage fallback rules used when source-record BIA or exposure fields are missing.

Add Security Zone Rule

Add Environment Rule

Add Business Service Rule

Add Application Service Rule

| <input type="checkbox"/> | Name (Scope)                           | Status | Exposure              | Business Impact      | Technical Impact | BOD Factors |
|--------------------------|----------------------------------------|--------|-----------------------|----------------------|------------------|-------------|
| <input type="checkbox"/> | DMZ<br>Security Zone                   | active | dmz / score 90        | Not set              | Not set          | Internet    |
| <input type="checkbox"/> | Management<br>Security Zone            | active | restricted / score 60 | Not set              | Not set          | Not set     |
| <input type="checkbox"/> | Production<br>Environment              | active | Not set               | high / score 80      | Not set          | Not set     |
| <input type="checkbox"/> | Lab<br>Environment                     | active | Not set               | low / score 10       | Not set          | Not set     |
| <input type="checkbox"/> | ERP<br>Business Service                | active | Not set               | critical / score 100 | Not set          | Not set     |
| <input type="checkbox"/> | Remote Access<br>Business Service      | active | Not set               | high / score 80      | Not set          | Not set     |
| <input type="checkbox"/> | Customer Portal<br>Application Service | active | internet / score 90   | medium / score 60    | Not set          | Internet    |

Figure 4 – The rule library provides a visible inventory of fallback context and its effect on exposure, impact, and CISA-style factors.

## 5. CISA-inspired decision path

Preserve the distinction between authoritative inputs and customer policy

CISA's published work provides two especially valuable principles for enterprise vulnerability prioritization. First, the Known Exploited Vulnerabilities catalog is an authoritative source of vulnerabilities exploited in the wild and should be an input to prioritization. Second, SSVC uses explicit decision points to guide vulnerability response based on exploitation and impact rather than relying on a single severity number.

C4SAM's BOD-style path evaluates four decision points: Publicly Exposed, In KEV, Automatable, and Technical Impact. These are the same four criteria used by CISA BOD 26-04. A finding is marked complete for this strict path only when a CVE exists and all required decision inputs are available. Missing CVE or missing enrichment causes the finding to enter the C4SAM True Risk fallback path. It still receives a priority and SLA, but it is not labelled BOD-complete or BOD-compliant.

| Decision point   | Primary source                               | Question                                                     | Typical values  |
|------------------|----------------------------------------------|--------------------------------------------------------------|-----------------|
| Publicly Exposed | C4SAM zone and reachability context          | Is the affected instance externally exposed?                 | Yes / No        |
| In KEV           | CISA KEV enrichment                          | Is there authoritative evidence of exploitation in the wild? | Yes / No        |
| Automatable      | CVE/SSVC/Vulnrichment or equivalent evidence | Can exploitation be performed reliably and at scale?         | Yes / No        |
| Technical Impact | CVE/SSVC/Vulnrichment                        | What control can the attacker gain?                          | Partial / Total |

### 5.1 Factor completeness is a control

A missing value is not the same as "No." C4SAM displays availability separately from the factor value. This prevents absent enrichment from being interpreted as negative evidence and makes data quality visible to analysts and auditors.

**Description**

The following response / ICMP packet has been received: - ICMP Type: 14 - ICMP Code: 0

The Timestamp Reply is an ICMP message which replies to a Timestamp message. It consists of the originating timestamp sent by the sender of the Timestamp as well as a receive timestamp and a transmit timestamp.

**What is the impact?:** This information could theoretically be used to exploit weak time-based random number generators in other services.

This vulnerability was found on Sun 12 Jan 2025 19:08.

Severity: **Low (2.1)**

CISA BOD 26-04: **True Risk fallback**

**BOD 26-04 factors**

| Factor           | Availability | Control       |
|------------------|--------------|---------------|
| CVE              | Missing      | Missing       |
| Publicly Exposed | Available    | Yes           |
| In KEV           | Available    | Yes           |
| Automatable      | Available    | Yes           |
| Technical Impact | Available    | Total control |

**Solution**

**Mitigation**

Various mitigations are possible: - Disable the support for ICMP timestamp on the remote host completely - Protect the remote host by a firewall, and block ICMP packets passing through the firewall in either direction (either completely or only for untrusted networks)

**True Risk / Remediation Priority**

Priority **P1** True Risk Score **48** **True Risk fallback**

| Threat Score | Exposure Score | Technical Impact Score | Business Impact Score |
|--------------|----------------|------------------------|-----------------------|
| 2.1          | 90             | 100                    | 50                    |

**Recommended SLA:** Immediate action / 3 days / forensic triage required if exploited or total-control risk

**CISA BOD 26-04 status:** Alternative True Risk fallback: no CVE available

**BOD 26-04 factors**

| Factor           | Availability | Control       |
|------------------|--------------|---------------|
| CVE              | Missing      | Missing       |
| Publicly Exposed | Available    | Yes           |
| In KEV           | Available    | Yes           |
| Automatable      | Available    | Yes           |
| Technical Impact | Available    | Total control |

**Tags:** CISA\_KEY, INTERNET\_FACING, PUBLICLY\_EXPOSED\_BOD\_26\_04, AUTOMATABLE\_EXPLOIT, TOTAL\_SYSTEM\_CONTROL, DMZ\_ASSET, NIS2\_RELEVANT\_SERVICE

Demo scoring: This PoC uses enriched sample data for Business Impact, Exposure, and Technical Impact. Production use requires validated CMDB, BIA, and exposure data.

Figure 5 – Vulnerability details show factor availability, strict/fallback status, component scores, reason codes, and the recommended operational SLA.

## 5.2 BOD-style timeline examples

The current C4SAM implementation includes configurable timeline defaults that follow the decision logic and remediation tiers of CISA BOD 26-04 (June 10, 2026): 3 days with forensic triage for the highest-risk combinations, 14 days for moderate risk, 60 days for lower risk, and fix on next system upgrade for the least critical findings. Outside the U.S. Federal Civilian Executive Branch, these timelines are customer-policy defaults. Verify the current official remediation timelines on the CISA directive page before making any compliance statement.

| Public exposure | KEV | Automatable | Technical impact | Example target                         |
|-----------------|-----|-------------|------------------|----------------------------------------|
| Yes             | Yes | Any         | Total control    | 3 days plus forensic triage            |
| Yes             | Yes | Yes         | Partial control  | 3 days                                 |
| Yes             | Yes | No          | Partial control  | 14 days                                |
| Yes             | No  | Yes         | Partial control  | 14 days                                |
| Yes             | No  | No          | Partial control  | 60 days                                |
| No              | Yes | Any         | Any              | 14 days                                |
| No              | No  | Yes         | Any              | 60 days                                |
| No trigger      | No  | No          | Any              | Fix on system upgrade or policy review |

Governance rule Keep external-source status, factor completeness, C4SAM priority, and customer SLA as separate fields. This prevents a configurable business policy from being misrepresented as an external compliance result.

## 6. From vulnerability to remediation

### Closing the security-to-operations execution gap

Prioritization creates value only when it changes execution. C4SAM connects the vulnerability decision to the affected Cockpit object, resolves that object to a configuration item in iTop, pre-populates remediation context, and creates a change request. The vulnerability view then displays the linked change reference and status at the specific affected-instance level.

### 6.1 End-to-end workflow

1. Ingest and normalize vulnerability findings from scanners and related evidence sources.
2. Resolve findings to affected assets, network devices, segments, zones, environments, applications, and business services.
3. Enrich CVE findings with KEV and available SSSVC/Vulnrichment decision data.
4. Calculate component scores, evaluate escalation rules, assign P1–P5, and record reason codes and factor completeness.
5. Select the exact affected Cockpit object and match it to the corresponding iTop configuration item.
6. Create a change request with priority, due date, affected object, risk rationale, and proposed solution.
7. Track the change reference and status in the vulnerability view without marking unrelated instances as remediated.
8. Verify remediation through rescanning or equivalent control evidence, then close or reassess residual risk.

### 6.2 Instance-level integrity

A grouped vulnerability may affect many systems with different exposures and priorities. A change request for one host must not imply that every instance is under remediation. C4SAM scopes local tracking to the vulnerability plus affected object and can display an object-specific row under a grouped finding. This preserves both portfolio readability and execution accuracy.

| Vulnerability                               | Change Request    | True Risk Priority | True Risk Score | Scope           | Context    | Reason Codes                               | Known Exploit | Vendor Fix | AV | AC | PR | UI | Reported time    | Last Reported    |
|---------------------------------------------|-------------------|--------------------|-----------------|-----------------|------------|--------------------------------------------|---------------|------------|----|----|----|----|------------------|------------------|
| ICMP Timestamp Reply Information Disclosure | No Change Request | P1                 | 48              | Public Zone     | Internet   | INTERNET_FACING PUBLICLY_EXPOSED_800_24_04 |               |            |    |    |    |    | 2025/01/12 19:08 | 2025/01/12 19:08 |
| ICMP Timestamp Reply Information Disclosure | C-000036          | P1                 | 48              | Public Zone     | Internet   | INTERNET_FACING PUBLICLY_EXPOSED_800_24_04 |               |            |    |    |    |    | 2025/01/12 19:08 | 2025/01/12 19:08 |
| ICMP Timestamp Reply Information Disclosure | No Change Request | P4                 | 41              | Restricted Zone | restricted | AUTOMATABLE_EXPLOIT REMOTE_CODE_EXECUTION  |               |            |    |    |    |    | 2025/01/12 19:08 | 2025/01/12 19:08 |
| ICMP Timestamp Reply Information Disclosure | No Change Request | P5                 | 35              | Work Zone       | Internal   | AUTOMATABLE_EXPLOIT REMOTE_CODE_EXECUTION  |               |            |    |    |    |    | 2025/01/12 19:08 | 2025/01/12 19:08 |
| ICMP Timestamp Reply Information Disclosure | No Change Request | P5                 | 35              | Guest Zone      | Internal   | AUTOMATABLE_EXPLOIT REMOTE_CODE_EXECUTION  |               |            |    |    |    |    | 2025/01/12 19:08 | 2025/01/12 19:08 |
| ICMP Timestamp Reply Information Disclosure | No Change Request | P5                 | 35              | SCADA LB        | Internal   | AUTOMATABLE_EXPLOIT REMOTE_CODE_EXECUTION  |               |            |    |    |    |    | 2025/01/12 19:08 | 2025/01/12 19:08 |

Figure 6 – One vulnerability appears across multiple contexts; only the specific public-zone instance carries the linked change reference.

### 6.3 Change creation and CI matching

The change workflow presents affected Cockpit objects and candidate iTop CIs, including match confidence. The user reviews the selected object, linked CI, priority, due date, vulnerability description, score, and solution before creation. This human confirmation is important where names, identifiers, or synchronization boundaries differ between systems.

Apply to

Hosts

Affected Cockpit objects

minion-demo.c4sam.com (computer) ×

iTop CI match 9 synced classes

minion-demo.c4sam.com (NetworkDevice) - direct strong

Manual iTop search

minion-demo.c4sam.com

**Remediate P1 vulnerability on minion-demo.c4sam.com**  
ICMP Timestamp Reply Information Disclosure

Affected Cockpit object  
**minion-demo.c4sam.com**

Linked iTop CI  
**minion-demo.c4sam.com** direct strong

P1 Due 15/07/2026

Vulnerability: ICMP Timestamp Reply Information Disclosure  
True Risk Priority: P1  
True Risk Score: 48  
Affected Cockpit objects: minion-demo.c4sam.com  
Solution: Various mitigations are possible:

- Disable the support for ICMP timestamp on the remote host completely
- Protect the remote host by a firewall, and block ICMP packets passing through the firewall in either direction (either completely or only for untrusted networks)

Search selected object in iTop

Close Create in iTop

Figure 7 – C4SAM prepares a change request for the selected affected object and matched iTop CI, with the risk decision and remediation text carried into execution.

**C-000036**  
● New (Routine Change)

Properties Cls (1) Contacts Work orders Related requests Related incidents Related problems

**General Information**

Ref C-000036

Organization C4SAM

Status **New**

Title Remediate P1 vulnerability on minion-demo.c4sam.com

**Relations**

Parent change undefined

**Dates**

Creation date 2026-07-12 09:43:45

Last update 2026-07-12 09:43:45

**Description**

Vulnerability: ICMP Timestamp Reply Information  
Disclosure True Risk Priority: P1 True Risk Score: 48  
Affected Cockpit objects: minion-demo.c4sam.com  
Solution: Various mitigations are possible: - Disable the support for ICMP timestamp on the remote host completely - Protect the remote host by a firewall, and block ICMP packets passing through the firewall in either direction (either completely or only for untrusted networks)

**Contacts**

Caller AIOPS SVC

Figure 8 – The resulting iTop Routine Change retains the C4SAM priority, score, affected object, solution, caller, and creation timestamp.

## 6.4 Production integration principles

- Use a backend integration service for production credentials, authorization, validation, retries, and audit logging.
- Validate allowed iTop classes and resolve synchronized Cockpit identifiers to authoritative iTop FunctionalCI keys.
- Make ticket creation idempotent so retries do not create duplicate changes.
- Persist cross-system references in the backend rather than relying on browser-local storage.
- Synchronize change status and closure evidence back to the vulnerability record.
- Do not equate change closure with vulnerability remediation until technical verification succeeds.

**PoC boundary** The current demo can use an optional browser-side iTop REST fallback when the backend proxy is unavailable. This is explicitly unsuitable for production because credentials, validation, persistence, and audit controls belong on the server side.

## 7. The operating model

Shared accountability across security, service ownership, and IT operations

True Risk Priority is most effective when adopted as an operating model rather than deployed as a scoring feature. The process needs clear decision rights, service-level objectives, exception handling, and a regular forum where security risk and change capacity are managed together.

### 7.1 Roles and responsibilities

| Role                           | Accountability                                                                                                              |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| CISO / security leadership     | Own prioritization policy, threat inputs, escalation criteria, risk acceptance, and security reporting.                     |
| CIO / technology leadership    | Ensure remediation capacity, service ownership, CMDB quality, and alignment with lifecycle and resilience priorities.       |
| Vulnerability management       | Operate ingestion, enrichment, triage, reason-code review, data-quality checks, and verification.                           |
| IT operations / platform teams | Assess feasibility, plan and execute changes, document rollback, and provide implementation evidence.                       |
| Application and service owners | Validate business context, approve service risk, coordinate testing and downtime, and accept residual risk when authorized. |
| Change authority / CAB         | Apply urgency-appropriate approval paths without allowing governance to erase time-critical response.                       |
| CMDB / architecture owners     | Maintain identifiers, relationships, zones, publication paths, and authoritative configuration context.                     |
| Risk, compliance, and audit    | Review policy versions, exceptions, evidence quality, timeliness, and external-claim accuracy.                              |

### 7.2 Decision cadence

| Cadence                   | Purpose                                                                      | Typical participants            |
|---------------------------|------------------------------------------------------------------------------|---------------------------------|
| Continuous / event driven | Recalculate when KEV, exposure, asset relation, or change status changes     | Platform automation             |
| Daily                     | Review new P1/P2 findings, failed enrichment, overdue urgent changes         | VM, SOC, operations             |
| Weekly                    | Balance remediation demand with maintenance windows and service risk         | VM, operations, service owners  |
| Monthly                   | Review policy performance, exception aging, ownership gaps, and data quality | CISO, CIO delegates, risk, CMDB |
| Quarterly                 | Approve material model changes and assess business outcomes                  | CISO, CIO, governance board     |

### 7.3 Exception and risk-acceptance control

A missed deadline must not be hidden inside a generic exception. Each exception should record the affected instance, priority and reasons, business owner, technical owner, constraint, compensating controls, residual-risk

statement, approval authority, expiry date, and next review. Expired exceptions should automatically re-enter the active queue.

Emergency changes need proportional controls: rapid authorization, pre-defined rollback, evidence capture, post-implementation review, and incident linkage where exploitation is suspected. The goal is not to bypass change management, but to make change governance responsive to risk.

## 8. Governance and control framework

Make the model trustworthy, reproducible, and auditable

| Control objective       | Required practice                                                                     | Evidence                                                    |
|-------------------------|---------------------------------------------------------------------------------------|-------------------------------------------------------------|
| Model integrity         | Version weights, thresholds, escalation rules, and SLA labels; test before activation | Policy version, approver, scenario results, activation date |
| Source integrity        | Identify authoritative source and freshness requirement for each input                | Source registry, timestamps, failed-ingestion alerts        |
| Decision explainability | Retain component scores, reason codes, factor values, and completeness                | Decision record per affected instance                       |
| External-claim accuracy | Separate CISA inputs and strict status from C4SAM fallback and customer policy        | Field definitions, UI labels, report disclaimers            |
| Workflow integrity      | Link the exact affected object and authoritative CI; prevent duplicates               | Cross-system IDs, match confidence, idempotency key         |
| Closure integrity       | Require remediation verification after implementation                                 | Rescan result, version evidence, control validation         |
| Access control          | Restrict model changes, overrides, change creation, and risk acceptance               | Role assignments and audit log                              |
| Data protection         | Keep integration credentials server-side and minimize sensitive context               | Secrets management, API controls, logging policy            |

### 8.1 Overrides

Manual override is necessary for exceptional intelligence or incorrect source data, but it is a privileged decision. An override should capture the original result, replacement priority, reason, evidence, user, time, expiry, and approval. Expiring overrides force the model to reconsider the finding when conditions change.

### 8.2 Data freshness

True Risk is time-sensitive. KEV status can change; an asset can become internet reachable; a service can move into production; a change can fail; and ownership can be assigned. C4SAM should recalculate on material events and display the calculation timestamp and source freshness. Outdated context should be visible rather than silently trusted.

### 8.3 Model limitations

- A score is only as reliable as asset resolution and context quality.
- Absence from KEV is not proof that exploitation is impossible or inactive.
- Public-zone labels and internet reachability can disagree; direct evidence should prevail.
- Business-impact metadata can be subjective and needs owner validation.
- Fallback rules increase coverage but also increase inference risk.
- Priorities guide decisions; they do not replace incident response, threat hunting, or engineering judgment.

## 9. Implementation architecture

Integrate evidence, context, decisioning, and execution

| Layer        | Purpose                                                    | Typical systems/data                                                                  |
|--------------|------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Evidence     | Detect weaknesses and enrich threat/impact                 | Vulnerability scanners, CVE, CISA KEV, SSVC/Vulnrichment, vendor advisories           |
| Context      | Resolve where the finding exists and what depends on it    | C4SAM Cockpit, CMDB, assets, network zones, services, applications, environments, BIA |
| Decision     | Calculate score, apply escalation, assign priority and SLA | C4SAM True Risk model, rule library, policy versions, reason codes                    |
| Execution    | Create and manage controlled remediation work              | iTop change, work orders, maintenance windows, approvals, rollback plans              |
| Verification | Confirm risk reduction and update status                   | Rescan, configuration/version checks, control validation, exception review            |
| Governance   | Measure outcomes and manage policy                         | Dashboards, audit records, access control, risk acceptance, executive reporting       |

### 9.1 Minimum viable data contract

For each affected instance, the platform should be able to identify the vulnerability, affected object, authoritative CI, network/security context, environment and service context, relevant CVEs, KEV status, available exploitability and impact evidence, component scores, reason codes, priority, target date, owner, change reference, and verification state. Where a field is missing, the absence must be explicit.

### 9.2 Event-driven recalculation

Recalculation should occur when a finding is added or updated, KEV or enrichment changes, asset-to-service relationships change, exposure changes, BIA metadata changes, an override is applied or expires, or a change transitions state. Event-driven recalculation reduces the risk that a once-low-priority item remains hidden after becoming exposed or exploited.

### 9.3 Reliability and security requirements

- Use stable identifiers and controlled matching across scanners, Cockpit, CMDB, and iTop.
- Record provenance for every decisive input and protect against stale synchronization.
- Make integrations idempotent, observable, retryable, and rate-limit aware.
- Apply least privilege to read context, change policy, create tickets, approve changes, and accept risk.
- Avoid browser-resident production secrets; use managed credentials and server-side APIs.
- Retain an audit history of policy and decision changes without exposing sensitive vulnerability details broadly.

## 10. Implementation roadmap

Adopt the operating model in controlled stages

| Phase                    | Focus                                                                                   | Exit criteria                                               |
|--------------------------|-----------------------------------------------------------------------------------------|-------------------------------------------------------------|
| 0. Sponsorship and scope | Define business outcomes, services in scope, risk appetite, owners, and governance      | Named executive sponsors, approved scope, baseline metrics  |
| 1. Data foundation       | Resolve assets and CIs; validate zones, internet exposure, services, BIA, and ownership | Measured coverage and freshness; critical services mapped   |
| 2. Decision pilot        | Enable scoring and reason codes for a bounded portfolio; compare against current triage | Stakeholders understand and accept priority rationale       |
| 3. Change integration    | Create iTop changes from selected affected instances; synchronize status                | End-to-end traceability and duplicate prevention            |
| 4. Policy governance     | Approve thresholds, escalations, SLA labels, exceptions, and override controls          | Versioned policy and recurring governance cadence           |
| 5. Scale and optimize    | Expand coverage, automate recalculation, improve data quality, tune reporting           | Sustained SLA performance and measurable exposure reduction |

### 10.1 Pilot design

A good pilot is limited in scope but representative: one externally exposed service, one high-value internal service, one standard workplace or infrastructure segment, and one non-production environment. This reveals whether exposure, impact, and lifecycle differences produce sensible priorities. The pilot should run in parallel with the existing process long enough to compare decisions and uncover data defects.

Success does not mean the model repeats every past analyst decision. Success means that disagreements can be traced to evidence, policy, or missing context, and then resolved deliberately.

### 10.2 Adoption risks and mitigations

| Risk                                               | Mitigation                                                                                                               |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Teams perceive the model as another security score | Tie each priority to operational treatment, owner, due date, and change status.                                          |
| CMDB gaps undermine confidence                     | Publish coverage metrics, expose missing fields, and prioritize context for critical services.                           |
| Urgent queue remains too large                     | Review escalation specificity, duplicate findings, exposure evidence, and service scope; do not simply raise thresholds. |
| Operations cannot meet targets                     | Align targets with emergency/standard change paths, capacity planning, and compensating-control options.                 |
| External guidance is overstated                    | Use approved terminology and preserve the strict/fallback distinction in UI, reports, and customer communication.        |

# 11. Measurement and executive reporting

## Measure reduced exposure and accountable execution

Executive reporting should show whether the organization is reducing material exposure, not just processing tickets. Volume counts remain useful for capacity planning, but they should be accompanied by priority, context, age, ownership, change, and verification measures.

| Metric                         | What it reveals                            | Caution                                       |
|--------------------------------|--------------------------------------------|-----------------------------------------------|
| Open P1/P2 affected instances  | Current urgent exposure                    | Count instances, not only unique CVEs         |
| Time to triage / assign owner  | Decision and accountability speed          | Separate source-ingestion latency             |
| Time to approved change        | Security-to-operations handoff             | Segment emergency and standard paths          |
| Time to verified remediation   | Actual risk reduction                      | Do not use ticket closure alone               |
| SLA attainment by priority     | Policy execution                           | Exclude only approved, time-bound exceptions  |
| KEV exposure backlog           | Known exploited risk                       | Distinguish public and internal instances     |
| Publicly exposed P1 aging      | Highest-confidence external attack surface | Validate reachability freshness               |
| Context completeness           | Model reliability and CMDB/BIA maturity    | Report by critical service and source         |
| Ownership gaps                 | Blocked accountability                     | Measure age and business impact               |
| Change success / rollback rate | Operational quality                        | Do not incentivize avoidance of urgent change |
| Reopened after verification    | Remediation effectiveness                  | Investigate scanner and deployment quality    |

## 11.1 Executive dashboard narrative

- Current exposure: urgent findings by service, environment, security zone, and owner.
- Movement: new, remediated, accepted, and overdue risk since the previous period.
- Execution: changes planned, approved, in progress, failed, and awaiting verification.
- Data confidence: percentage with authoritative asset, exposure, BIA, service, CVE, and enrichment context.
- Constraints: capacity, unsupported platforms, maintenance dependencies, and recurring ownership gaps.
- Decisions required: risk acceptance, investment, service downtime, lifecycle acceleration, or control redesign.




Figure 9 – True Risk mode changes the operational view from scanner severity to instance-specific priority, score, scope, context, reason codes, and change status.

## 12. Illustrative decision scenarios

How the same model supports different operational outcomes

| Scenario                               | Evidence                                                     | Likely treatment                             | Management rationale                                                       |
|----------------------------------------|--------------------------------------------------------------|----------------------------------------------|----------------------------------------------------------------------------|
| Internet-facing KEV with total control | Public + internet, KEV, total control                        | P1; expedited change and forensic triage     | Active threat path with severe technical consequence                       |
| Critical ERP weakness, internal only   | High business impact, internal exposure, no exploit evidence | P2/P3 depending score and policy             | Business consequence is high, but reachability changes urgency             |
| Moderate public finding without CVE    | Public + internet, incomplete enrichment                     | P1 fallback with explicit missing-CVE status | Reachability and context justify action without a false BOD-complete claim |
| Automatable issue in restricted zone   | Internal/restricted, automatable, moderate impact            | P4 or P3 based on score and service context  | Automation matters, but exposure and impact constrain urgency              |
| Low-risk lab finding                   | Isolated lab, low business impact, no active exploit         | P5 / lifecycle                               | Routine remediation should not displace material exposure                  |
| Internet-facing asset with no owner    | Reachable, ownership gap                                     | At least P2 under configured policy          | Lack of ownership increases execution risk and requires management action  |

These scenarios are illustrations, not universal outcomes. The decisive requirement is that the output follows approved policy, shows its evidence, and leads to an accountable treatment path.

## 13. Recommendations for CIOs and CISOs

---

A joint agenda for reducing material cyber exposure

1. Adopt one shared priority language across vulnerability management, service ownership, and IT operations.
2. Treat internet reachability, KEV evidence, and technical control as decisive signals rather than optional report columns.
3. Invest in the context graph: authoritative assets, zones, services, applications, BIA, data classification, and ownership.
4. Require each urgent finding to have an affected instance, reason trail, owner, target, and controlled remediation path.
5. Integrate change creation and status so security priority becomes visible operational work.
6. Keep CISA-derived evidence, C4SAM fallback logic, customer policy, and compliance assertions explicitly separate.
7. Measure verified risk reduction, not only scans, findings, patches, or tickets.
8. Govern model changes like other material risk-policy changes: version, test, approve, monitor, and review.

Board-level outcome A mature True Risk program gives leadership a credible answer to three questions: Are the exposures most likely to cause material harm being addressed first? Is remediation owned and progressing through controlled change? Can the organization prove the decision and the result?

## A. Appendix: field and evidence model

| Domain           | Core fields                                                                              | Quality expectation                             |
|------------------|------------------------------------------------------------------------------------------|-------------------------------------------------|
| Vulnerability    | Finding ID, name, scanner severity, score, CVEs, solution, first/last seen               | Normalized, deduplicated, timestamped           |
| Threat           | KEV, exploit evidence, public exploit, ransomware indicator                              | Source and retrieval time retained              |
| Exposure         | Zone, exposure level/score, internet facing, publicly exposed, publication path          | Direct evidence distinguished from inference    |
| Technical impact | Automatable, RCE, partial/total control                                                  | Value plus availability and provenance          |
| Business impact  | Environment, BIA score, criticality, RTO/RPO, service/application, data/regulatory scope | Owner-approved and reviewed                     |
| Ownership        | Technical owner, service owner, business owner                                           | Accountable person or group, not free text only |
| Decision         | Component scores, total score, escalation, priority, reason codes, model version         | Reproducible and immutable history              |
| Treatment        | SLA, due date, exception, compensating control, acceptance                               | Approved authority and expiry                   |
| Change           | Cockpit object, iTop CI, change ref/status, implementation and rollback                  | Exact-instance linkage and idempotency          |
| Verification     | Rescan, version/control evidence, closure state                                          | Independent of ticket closure                   |

## B. Appendix: policy checklist

---

- The purpose, scope, executive owner, and risk appetite are documented.
- Threat, exposure, technical-impact, and business-impact sources are named and monitored.
- Weights total 100%, thresholds are approved, and escalation rules are explicit.
- Priority bands map to operational treatment and change paths.
- CISA-style factor completeness is separate from C4SAM fallback scoring.
- Missing data is visible and never silently interpreted as a negative value.
- Overrides and exceptions require reason, evidence, approval, and expiry.
- Affected-instance and CI matching are validated before change creation.
- Production integration credentials and validation remain server-side.
- Change closure requires technical verification before risk closure.
- Metrics measure exposure reduction, timeliness, data confidence, and change quality.
- The policy is reviewed at least quarterly and after material incidents or source changes.

## C. References and source notes

---

### CISA, BOD 26-04: Prioritizing Security Updates Based on Risk

Official directive (June 10, 2026) that sets risk-based remediation timelines using four criteria: public exposure, KEV status, exploit automation, and technical impact. Supersedes BOD 19-02 and BOD 22-01.

### CISA, Known Exploited Vulnerabilities Catalog

Authoritative catalog of vulnerabilities exploited in the wild; CISA recommends using it as an input to prioritization.

### CISA, BOD 22-01: Reducing the Significant Risk of Known Exploited Vulnerabilities

Directive establishing KEV remediation requirements for U.S. Federal Civilian Executive Branch agencies. Superseded by BOD 26-04 in June 2026; retained here as a historical reference.

### CISA, BOD 23-01: Improving Asset Visibility and Vulnerability Detection on Federal Networks

Establishes asset visibility as a precondition for effective vulnerability management.

### CISA, Stakeholder-Specific Vulnerability Categorization Guide

Decision-tree guidance for prioritizing vulnerability response based on exploitation and impact.

### C4SAM True Risk implementation

July 2026 implementation: True Risk model, BOD-style timeline logic, rule workbench, instance context, factor completeness, and vulnerability-to-iTop change workflow.

### C4SAM customer explanation deck

C4SAM True Risk Priority — How we calculate risk score, remediation priority, and BOD-style timelines, July 2026.

Source qualification CISA published BOD 26-04, “Prioritizing Security Updates Based on Risk,” on June 10, 2026. This whitepaper describes the C4SAM implementation as aligned with the BOD 26-04 decision criteria. Alignment is a design statement, not a certification: BOD 26-04 is binding only for U.S. Federal Civilian Executive Branch agencies, and CISA does not certify commercial products.

## D. Appendix: glossary of terms and abbreviations

| Term         | Meaning                                                                                                  |
|--------------|----------------------------------------------------------------------------------------------------------|
| BIA          | Business Impact Analysis; assessment of how critical a service or system is for the organization.        |
| BOD          | Binding Operational Directive; a compulsory CISA directive for U.S. federal civilian agencies.           |
| CAB          | Change Advisory Board; the body that reviews and approves changes.                                       |
| CI           | Configuration Item; an object managed in a CMDB, such as a server or an application.                     |
| CISA         | U.S. Cybersecurity and Infrastructure Security Agency.                                                   |
| CMDB         | Configuration Management Database; the inventory of IT assets and their relationships.                   |
| Cockpit      | The C4SAM module that models assets, zones, services, and their relationships.                           |
| CVE          | Common Vulnerabilities and Exposures; standard identifier for a publicly known vulnerability.            |
| CVSS         | Common Vulnerability Scoring System; standard technical severity score for vulnerabilities.              |
| DMZ          | Demilitarized Zone; a network segment between the internet and internal networks.                        |
| ERP          | Enterprise Resource Planning; central business software, for example for finance and logistics.          |
| iTop         | Open-source IT service management and CMDB tool used here for change requests.                           |
| KEV          | Known Exploited Vulnerabilities; the CISA catalog of vulnerabilities exploited in the wild.              |
| PoC          | Proof of Concept; an early implementation used to demonstrate feasibility.                               |
| RCE          | Remote Code Execution; the ability of an attacker to run code on a system remotely.                      |
| RTO / RPO    | Recovery Time Objective / Recovery Point Objective; maximum acceptable downtime and data loss.           |
| SLA          | Service Level Agreement; an agreed target, for example a remediation deadline.                           |
| SOC          | Security Operations Center; the team that monitors and responds to security events.                      |
| SSVC         | Stakeholder-Specific Vulnerability Categorization; CISA decision-tree method for vulnerability response. |
| Vulnrichment | CISA program that enriches CVE records with SSVC decision data.                                          |

## About C4SAM

---

C4SAM — Code Name: Cockpit for Situation Awareness and Monitoring — helps organizations understand and secure their digital infrastructure by connecting operational visibility, vulnerability information, business and service context, and controlled response. The True Risk Priority capability extends that mission from finding vulnerabilities to deciding and executing what matters first.

The practical promise One prioritized view. One reason trail. One accountable remediation workflow.  
Evidence from detection through verified risk reduction.