CSURFACE

WHITEPAPER · THREAT INTELLIGENCE

Vulnerability prioritization beyond CVSS

Threat Intelligence and dynamic prioritization: how to get the remediation queue right by combining technical severity, real exploitability (exploitation probability index and active-exploitation catalog) and asset criticality.

CAPABILITY	AUDIENCE	EDITION
Threat Intelligence	CISO · Vuln. Manager · SOC	2026

READING	CLASSIFICATION	SOURCE
9 pages	Public	CSURFACE Platform

EXECUTIVE SUMMARY

The right remediation queue is dynamic

The thesis of this whitepaper, in one page — what executives need to know before the technical detail.

CVES BEHIND BREACHES

5%

of CVEs account for 20% of actual breaches

WEAPONIZATION WINDOW

24h

28% of exploited vulnerabilities are weaponized within 24 hours

TYPICAL MONTHLY REMEDIATION

15%

of the backlog is fixed per month — insufficient to keep pace with the rate of intake

Most Vulnerability Management programs prioritize by CVSS. It is an intuitive and technically insufficient choice. CVSS quantifies the **theoretical technical severity** of a vulnerability — the impact it would produce *if* it were exploited. It does not indicate whether it **is being exploited in the present**, nor does it weigh the criticality of the asset on which it resides.

The result is a queue systematically misaligned with real risk. Teams allocate scarce capacity to vulnerabilities with high CVSS and low exploitability, while medium-severity vulnerabilities — yet actively exploited in ongoing campaigns — remain open. Only **5% of CVEs account for 20% of breaches** (Verizon DBIR 2024); prioritizing by CVSS is, in practice, ignoring this concentration of risk.

This whitepaper presents CSURFACE's approach: **dynamic prioritization on three axes** — technical severity, real exploitability (via exploitation probability index and active-exploitation catalog) and asset criticality —, continuously reassessed as the threat landscape evolves. A CVE scored 7.5 can become the top priority the day after it enters the active-exploitation catalog. The queue must reflect this dynamic.

The thesis. Effective prioritization is a threat intelligence problem. The proper queue constitutes a dynamic ranking, reordered each time an exploit becomes available, a ransomware family incorporates the vulnerability or active-exploitation catalogs confirm it. Lists fixed at the moment of *disclosure* fall out of alignment within days.

What this whitepaper covers

- Why the volume of CVEs has grown — and why CVSS does not scale with it.
- The three reasons why CVSS, on its own, leads to the wrong queue.
- CSURFACE's three-axis prioritization model.
- How the intelligence feed stays dynamic and a real use case.

THE LANDSCAPE · PART 01

The growth of vulnerabilities — and the queue that does not recede

Publication volume, remediation capacity and the reason the central question has ceased to be quantitative.

~30 thousand
CVEs published in 2024 — 15% more than in 2023; 2025 approached 48 thousand

~15 thousand
vulnerabilities per year in the inventory of a mid-sized organization

15% / month
typical backlog remediation rate — below the rate of intake of new CVEs

Volume grew faster than capacity

The number of CVEs published every year has ceased to be an administrative detail and has become a capacity problem. There were around **30 thousand CVEs in 2024** — 15% more than in 2023 — and 2025 approached **48 thousand**, a jump on the order of 60%.

A mid-sized organization must handle something close to **15 thousand vulnerabilities per year** in its inventory. The typical remediation rate, however, stands at around **15% of the backlog per month** (Cyentia). The difference between intake and output is structurally negative: the backlog grows consistently.

Faced with this asymmetry, the central question ceases to be *"how many vulnerabilities do we have?"* and becomes: **"are we fixing the ones that represent real risk?"**

Key statistics

Breaches concentrated in few CVEs	5%→20%
Exploited and weaponized within 24h of disclosure	28%
Remediation typical monthly of the backlog	15%
Time average between disclosure and patch	20 days

Sources: Verizon DBIR 2024, CISA 2024, Cyentia, Snyk — public reports.

The cost of getting the queue wrong. The average cost of a data breach in the United States reached **$10.22M** in 2025 (IBM Cost of a Data Breach 2025). Every actively exploited vulnerability postponed in the queue represents direct exposure for the duration of the delay. Prioritizing with precision represents the distinction between a remediation program that effectively reduces risk and one that absorbs capacity without directing effort to the point of greatest probability of impact.

Why this matters. Remediation capacity is finite and will not keep up with the volume of CVEs. The only real lever is the order of the queue: fix first what has the greatest probability of being exploited, on the asset that matters most.

WHY CVSS FALLS SHORT · PART 02

CVSS measures theoretical impact — not real risk

The three reasons why prioritizing exclusively by the CVSS score leads to the wrong target.

BY CVSS · FIX FIRST

9.8

CVSS "critical" — yet with no public exploit, absent from ransomware kits and requiring the rarest of environmental conditions to work.

→

BY REALITY · THIS IS THE URGENCY

7.5

CVSS "high" — with a published public exploit, in the active-exploitation catalog and actively used by ransomware families in real campaigns.

One number, three blind spots

CVSS quantifies the theoretical impact *if* the vulnerability is exploited. It does not answer the operational question — is it **being exploited now?**. The contrast above sums up the problem: by CVSS, the team fixes the 9.8 first; by the reality of risk, the 7.5 is the urgency. The technical score, on its own, leads to the inverse of the correct decision.

The reason is structural. The three reasons that follow explain why prioritizing exclusively by the CVSS score leads to the wrong target — and why none of them is resolved by calibrating CVSS itself more carefully.

Reason 01 · CVSS measures severity, not exploitability. The score answers "how serious it would be", not "this is being exploited". It is the axis the example above exposes.

Reason 02 · CVSS ignores asset context. A 9.8 vulnerability on a production server with cardholder data is catastrophic. The same 9.8 on an isolated test machine, with no data, is practically irrelevant. CVSS assigns both the same number — and any queue built solely on it treats incomparable cases as equal.

WHAT CSURFACE OBSERVES · THE PLATFORM'S OWN DATA

5.7%
of the 2,361 vulnerabilities analyzed have a public exploit — and only 1.8% are in the active-exploitation catalog

19%
of a threat feed with an average CVSS of 9.0 require immediate action under SSVC categorization

~5 days
is the average time until a public exploit emerges after a flaw is disclosed

Reason 03. CVSS is static; risk changes continuously

The score is assigned at *disclosure* and rarely revised. The threat landscape, however, evolves within days. Consider the typical trajectory of a single CVE:

1. **Day 0 — disclosure.** The CVE receives CVSS 7.5. The queue places it as medium priority.
2. **Day 3 — proof of concept.** A functional exploit is published in a public repository. The barrier to attack collapses.
3. **Day 7 — criminal adoption.** Ransomware families begin incorporating the vulnerability into their campaigns.
4. **Day 14 — official confirmation.** The CVE enters the active-exploitation catalog, attesting to active exploitation in real attacks.

CVSS remains 7.5 from day 0 to day 14. Real risk, in that interval, rose several times over — and the static queue never registered the change.

THE CSURFACE APPROACH · PART 03

Dynamic prioritization on three axes

Technical severity, real exploitability and asset criticality — combined into a single living score.

THE THREE AXES OF DYNAMIC PRIORITIZATION

1

Technical severity

CVSS as a *baseline* of impact, on a scale of 0 to 10. A stable axis — it rarely changes after disclosure.

2

Real exploitability

The concrete probability of exploitation, informed by exploitation probability index, active-exploitation catalog, active intelligence and ransomware.

3

Asset criticality

The weight of the asset adjusts the position of the same vulnerability according to where it appears in the organization.

What makes up real exploitability

Axis 02 is what CVSS does not cover. It brings together independent pieces of exploitation evidence, each with its own update cadence:

CRITICAL **Exploitation probability**
Estimates the statistical probability of exploitation in the next 30 days, updated daily.

CRITICAL **Active-exploitation catalog**
Catalog of vulnerabilities demonstrably exploited in real attacks. Being listed in it means automatic priority.

ATTENTION **Active intelligence**
Monitoring of exploitation kits, criminal trade forums and public proof-of-concept repositories.

ATTENTION **Ransomware indicators**
Association with active families — a signal of operational adoption of the vulnerability by attackers.

The final score

The three axes combine into a single priority reading, continuously reassessed:

> The platform prioritizes by observed active exploitation and asset criticality, combining technical severity and threat intelligence — in an auditable way, without requiring the customer to know the internal weighting mechanics.

Observed active exploitation on a critical asset takes the top; a high severity with no signal of exploitation recedes in the queue.

How the asset is classified. CSURFACE assesses data sensitivity (personal and sensitive data under the CCPA and state privacy laws, cardholder data under PCI), exposure (public, intranet or segregated), business criticality and position in the chain — from customer-facing frontend to backoffice.

Why exploitation weighs. A high severity with no observed exploitation recedes in the queue, while active exploitation on a critical asset takes the top. The queue reflects the prevailing risk, and not merely the theoretical impact of the vulnerability.

THE CSURFACE APPROACH · PART 03

An intelligence feed continuously updated

The integrated sources and the closed loop that reorders the priority queue without dependence on manual intervention.

THE CLOSED LOOP OF REASSESSMENT — WITHOUT MANUAL INTERVENTION

1. **Automatic reassessment**
   The priority score is recalculated without any manual trigger when the CVE changes status.
2. **Targeted notification**
   Material changes — such as the entry of a CVE into the active-exploitation catalog — generate an immediate alert.
3. **Synchronized ticket**
   The ticket in Jira or ServiceNow is updated, and the SLA tracks the new priority.
4. **Executive dashboard updated**
   The management dashboard reflects the reordered queue in real time, dispensing with the periodic reporting cycle.

Integrated sources

Prioritization is only dynamic because the intelligence feed is too. CSURFACE integrates and cross-references, continuously:

- CRITICAL **Exploitation probability**
  Daily exploitation probability feed.
- CRITICAL **Active-exploitation catalog**
  Active-exploitation catalog, updated within hours.
- **NVD**
  Base data of CVEs and their official metadata.
- **Vendor advisories**
  Security bulletins from the affected vendors.
- ATTENTION **Public proofs of concept**
  Including Exploit-DB and open repositories.
- ATTENTION **Criminal monitoring**
  Discussion and trade of exploits in closed forums.
- ATTENTION **Honeypot signals**
  When a vulnerability begins to be scanned at scale.

Why the feed must be continuous

A static queue would only suffice if the threat landscape were static too. An exploit can emerge three days after disclosure; active-exploitation catalogs can confirm the exploitation a few days later. Each of these changes alters the real priority of a CVE — and, if the feed does not register them on the same day, the queue operates on outdated data, disconnected from the prevailing risk.

That is why CSURFACE treats intelligence as a continuous flow. Each source alongside has its own cadence, and the closed loop described above ensures that every material change reaches the queue — and the responsible team — without depending on the individual perception of an analyst.

The cadence of each axis. Technical severity is practically static. Exploitability is reassessed daily — the exploitation probability index runs every day, the active-exploitation catalog within hours, active intelligence in real time. Asset criticality is reviewed monthly or with each configuration change.

APPLICATION AND RESULTS · PART 04

The same team, a different queue

A representative use case: how dynamic prioritization reordered the top of the queue without increasing capacity.

PRIORITIZATION BY CVSS

3

material incidents in 12 months — all in CVEs scored CVSS 7.x, yet present in the active-exploitation catalog.

→

CSURFACE DYNAMIC PRIORITIZATION

0

material incidents in the following 6 months, with the same team and the same remediation capacity.

Use case · mid-sized bank

A mid-sized financial institution operated with **4,200 assets** and around **8,700 active CVEs**. The team — four remediation staff and two developers — had capacity for about 80 fixes per month, against a backlog growing by 50 per month. The queue prioritized the 200 vulnerabilities scored CVSS 9.0 or higher.

The three incidents recorded in twelve months all occurred in vulnerabilities with CVSS in the 7.x range — below the "critical" cutoff — yet present in the active-exploitation catalog. The CVSS queue had left them out.

With dynamic prioritization, the top of the queue was reordered without changing the team's capacity.

How the top of the queue changed

47 vulnerabilities left the critical list
High CVSS, yet with very low exploitation probability and without any signal of exploitation.

CRITICAL **89 entered the top of the queue**
Medium CVSS, yet with elevated exploitation probability or already in the active-exploitation catalog.

ATTENTION **31 moved up in position**
Exclusively through the asset criticality.

The result. Remediation capacity came to be directed to where risk was effectively concentrated. In the following six months, material incidents fell from three to zero — without expanding the team or additional budget.

Capacity did not change — allocation did. The team maintained the pace of about 80 fixes per month. What changed was the composition of those fixes: effort was no longer absorbed by high-CVSS vulnerabilities without active exploitation and came to cover those that presented real exploitation on the assets of highest criticality.

APPLICATION AND RESULTS · PART 04

Metrics of change and integration in the platform

Variations in the program's indicators and how dynamic prioritization articulates with the other capabilities of the platform.

Metrics of change · mid-sized bank

INDICATOR	BEFORE · PRIORITIZATION BY CVSS	AFTER · DYNAMIC PRIORITIZATION	VARIATION
MTTR of critical exploited vulnerabilities	87 days	**18 days**	-79%
Fixes applied to CVEs with real exploitation	45%	**92%**	+47 pp
Monthly backlog balance	+50	**-12**	reversed
False positives in the critical queue	High	**Low**	Material reduction
Material incidents per year	3	**0** (in 6 months)	-100%

Indicators observed in the use case described on the previous page. Results vary according to the size of the inventory, remediation capacity and the maturity of each organization's program.

-79%
in the MTTR of critical exploited vulnerabilities — from 87 to 18 days

backlog
monthly balance reversed: from +50 vulnerabilities per month to -12

3 → 0
material incidents — from three in twelve months to zero in the following six months

Integration with the other capabilities

Threat Intelligence does not operate in isolation. On CSURFACE's unified platform, it consumes the other capabilities and feeds back into each of them:

Continuous discovery
Every new asset on the external surface automatically enters the scope of prioritization.

Exploitability validation
Active exploitability validation where test coverage exists, with passive monitoring across the rest of the surface.

CRITICAL **Leaked credential monitoring**
Compromised credentials combined with a critical vulnerability raise the priority to the maximum.

From queue to financial value. The calculated priority feeds CSURFACE's risk calculator, which estimates the expected annual loss by sector. The higher the priority of a vulnerability, the greater the cost of keeping it open — a reading that translates the technical queue into business language accessible to the board and the audit committee.

Key program indicators

- MTTR of the vulnerabilities effectively exploited — not of the overall backlog average.
- Percentage of remediation effort applied to CVEs with confirmed exploitation.
- Net monthly backlog balance: a trend toward reduction or growth.

NEXT STEPS

The right queue begins with threat intelligence

Threat Intelligence with dynamic prioritization is one of the capabilities of the CSURFACE platform — a Continuous Exposure Management platform. In a unified architecture, it brings together continuous discovery of the external surface with Machine Learning, analysis of the digital supplier chain, exploitability validation, threat intelligence with dynamic prioritization and leaked credential monitoring. It operates entirely externally — without an agent and without installation —, with optional cloud, WAF and CIEM integrations for organizations that require deeper visibility into internal environments.

Assess how your queue would reorder

In a 30-minute demonstration, we show how the current priority list would reorganize with three-axis dynamic prioritization.

Financial risk calculator

Estimate the expected annual loss by sector and quantify the cost of keeping each vulnerability open.

Whitepaper: ML Discovery

Continuous discovery of the external surface — the inventory base on which all dynamic prioritization must operate.

Discover your reordered queue — by real exploitation, not by the theoretical score

Request a preliminary analysis