CSURFACE

WHITEPAPER · EXPOSURE WINDOW

The exposure window between the scheduled scan and remediation

In three years, the average time between the publication of a vulnerability and the first exploit dropped from 32 to 5 days — and the cycle of traditional vulnerability management, point-in-time and monthly, no longer keeps pace.

CAPABILITY	AUDIENCE	EDITION
Continuous Exposure Management	CISO · Risk · VulnMgmt · SOC	2026

READING	CLASSIFICATION	SOURCE
14 pages	Public	CSURFACE Platform

EXECUTIVE SUMMARY

The exploit is faster — and your window is open

The thesis of this whitepaper on a single page, with data from CSURFACE's emerging threats telemetry: the speed of exploitation and the cycle of vulnerability management no longer meet.

Whitepaper 2026
CSURFACE

EXPLOITS IN ≤ 5 DAYS · 2024 VS. 2020

+183%

153 CVEs exploited in ≤ 5 days in 2024 — against 54 in 2020[1]

TYPICAL VM CYCLE

30 days

monthly scan — the minimum interval between snapshots

ZERO-DAYS BEFORE DISCLOSURE · 2026*

52%

of the zero-days observed in 2026 were already exploited before the CVE was published — against 21% in 2025[1]

CSURFACE's emerging threats telemetry records that **153 CVEs were exploited in ≤ 5 days after disclosure in 2024** — nearly three times the 54 of 2020. In 2026, in the first months alone, there are already 81. The jump occurred in 2024 and the annual band settled at a structurally higher level than in any previous year. The acceleration is what the telemetry has recorded year after year.

Traditional vulnerability management, based on periodic scans, operates at a pace inherited from another era. A monthly cycle — already considered good at many organizations — delivers a snapshot of the estate every 30 days. The quarterly or annual cycle delivers an even more outdated snapshot. Between one snapshot and the next, the exposure window remains open — and the attacker has enough time to get in.

This whitepaper analyzes the acceleration of exploitation over the last three years, sets out the mathematical gap between the attacker's speed and the traditional cycle, and presents the CSURFACE approach — continuous discovery, threat sensor and dynamic prioritization — to close that window in hours, not weeks.

The thesis. Exposure window is the interval between the moment a vulnerability becomes exploitable and the moment the organization applies the corresponding control. It can only be closed by an instrument that operates continuously, at the same cadence as the attacker.

What this whitepaper covers

- The acceleration of exploitation — a 3-year trend, with public data.
- Seven recent public cases — exploitation in hours and days.
- The anatomy of the exposure window and why the point-in-time cycle does not close it.
- The CSURFACE approach — threat sensor, continuous discovery and prioritization.

THE ACCELERATION · PART 01

From 54 to 153 in four years · the fast-exploit jump

The number of CVEs exploited in ≤ 5 days after disclosure nearly tripled between 2020 and 2024 — a reading from CSURFACE's own emerging threats sensor.

Absolute number of CVEs exploited in ≤ 5 days · per year

CSURFACE's emerging threats telemetry tracks, year by year, the absolute number of CVEs whose exploitation in production was observed within 5 days after public disclosure. The seven-year series shows a structurally new regime starting in 2024 — nearly three times the 2020 baseline:

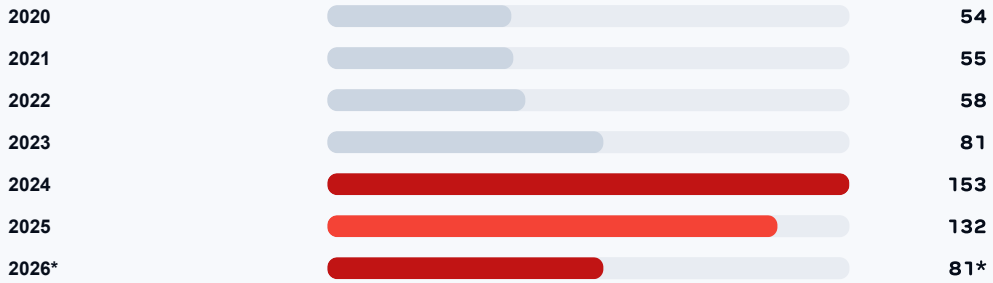

Source: CSURFACE emerging threats telemetry. *2026 covers only the first months of the year. The jump between 2023 (81) and 2024 (153) marks the entry into a structurally different regime — corroborated by Mandiant M-Trends 2024, which reports a drop in the median TTE to 5 days.[2]

What CSURFACE observed over the last three years

1. **Rising CVE volume.** The NVD published 18,362 CVEs in 2020. In 2025, there were **48,126** — nearly 2.6× the volume of five years earlier. The pressure on remediation queues increased in proportion.
2. **Zero-days doubled in 2024.** CSURFACE recorded 57 zero-days in 2023, **103 in 2024** and 89 in 2025. In 2026, there are already 64 in the first months — a pace that projects a new annual record.
3. **Exploitation before disclosure.** In 2025, 21% of zero-days were exploited before the CVE became public. In 2026, **52%** — more than half. The attacker arrives before the record even exists.
4. **The 2024 jump was structural.** The number of exploits in ≤ 5 days went from 81 (2023) to 153 (2024) — growth of 88% in one year. 2025 recorded 132; 2026, in a few months, already reaches 81. The annual band settled at a structurally new level.

The implication. The shift from 54-81 to 132-153 fast exploits per year is structural. It reflects a reorganization of the offensive ecosystem — automation, AI-assisted exploitation, specialization and an exploit market — that will endure. The defensive instrument must operate on a comparable time scale.

DYNAMIC EXPLOITATION · PART 02

Dormant CVEs that awoke · what CVSS alone ignores

Five real vulnerabilities whose exploitation probability more than doubled in the last weeks — observed by CSURFACE's emerging threats telemetry.

Whitepaper 2026
CSURFACE

Why this matters. CVSS expresses the *theoretical* severity of the flaw — assigned at publication and rarely updated. The **exploitation probability index** expresses the real probability of exploitation over the next 30 days and changes daily. A vulnerability can lie dormant for months with a low exploitation probability and, in a few days, have the index multiplied by 10× or 50× when exploits, PoCs or campaigns appear. Whoever prioritizes by CVSS alone misses that movement.

EXPLOITATION PROBABILITY · 30 DAYS · 5 CVES WITH HIGHEST GROWTH

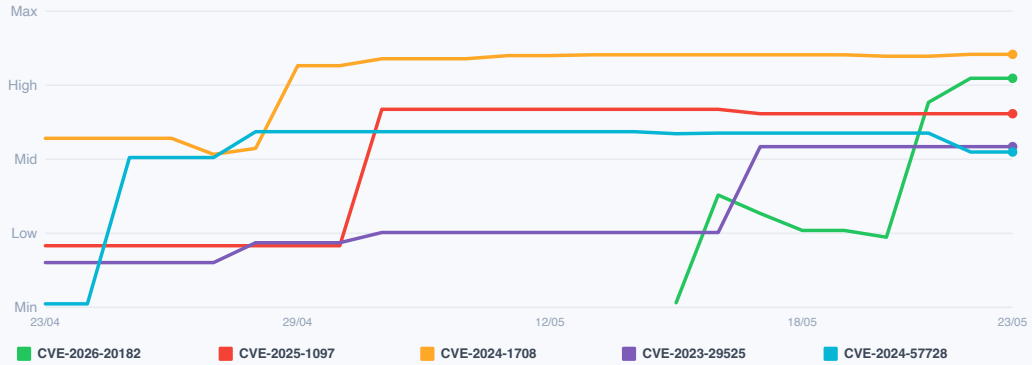

The five vulnerabilities in the chart

CVE	AFFECTED PRODUCT	PROBABILITY MIN → CURRENT	SIGNALS
CVE-2026-20182	Cisco Catalyst SD-WAN Manager	Low → Imminent	ACTIVE EXPLOIT
CVE-2025-1097	Kubernetes ingress-nginx	Moderate → High	rising
CVE-2024-1708	ConnectWise ScreenConnect	High → Critical	ACTIVE EXPLOIT
CVE-2023-29525	XWiki Platform	Low → High	dormant in 2023, reactivated
CVE-2024-57728	SimpleHelp Remote Access	Low → High	ACTIVE

Source: CSURFACE emerging threats telemetry. Daily capture of the exploitation probability index, the active exploitation catalog and public exploit availability. Snapshot: May 23, 2026.

What CVSS alone does not see

CVE-2026-20182 (Cisco Catalyst SD-WAN Manager) had a low exploitation probability on May 15. Seven days later, the index already ranked it among those with imminent exploitation — with no change in CVSS. A queue based on CVSS alone, recalculated once a month, would keep this CVE in the middle of the block. A dynamic queue pushes it to the top on the day the signal appears.

Intelligent and continuous prioritization. CSURFACE composes the queue by combining the current exploitation probability index, the 30-day delta, the active exploitation catalog, public exploit and active-exploitation telemetry — reflecting the present of the threat rather than a score assigned months earlier.

PUBLIC CASES · PART 02

Seven recent incidents · exploitation in hours and days

Seven public incidents with mass exploitation within 15 days after disclosure — all post-CVE, all within the scan cycle of most organizations.

The pattern behind the seven cases. Each incident below was exploited in production within a window of up to **15 days after public disclosure** — in several cases, in **less than 24 hours**. All post-CVE, all within the interval of a monthly scan. None of them would be caught in time by traditional vulnerability management.

CRITICAL **Log4Shell · CVE-2021-44228**

Disclosed on December 9, 2021. Cloudflare and Cisco Talos recorded large-scale exploitation attempts in the first hours after disclosure. Unauthenticated remote RCE in millions of Java applications. It continued to be exploited in production for months after the patch.

Time to mass exploitation · ≤ 24 hours

CRITICAL **Spring4Shell · CVE-2022-22965**

Disclosed on March 30, 2022. Public PoC in 24 hours, active exploitation in ~48 hours after disclosure. Mass exploitation by Mirai-variant botnets observed by multiple vendors. RCE in Spring Framework — a web stack widely adopted in enterprises.

Time to mass exploitation · 24-48 hours

CRITICAL **VMware vRealize · CVE-2022-22954**

Disclosed on April 6, 2022. Sophos, Rapid7 and Morphisec detected active exploitation for ~3 days after disclosure, with cryptominers and reverse shells. RCE via template injection in Workspace ONE Access — a target of lateral-movement operations in corporate networks.

Time to mass exploitation · ~3 days

CRITICAL **F5 BIG-IP · CVE-2022-1388**

Disclosed on May 4, 2022. Public PoCs in ~4 days, active mass exploitation around May 9. An authentication bypass in iControl REST allowed RCE with root privilege. Hundreds of exposed instances compromised within a few days.

Time to mass exploitation · ~5 days

CRITICAL **Citrix Bleed · CVE-2023-4966**

Disclosed on October 10, 2023. Post-disclosure mass exploitation observed in ~10 to 14 days — affecting Boeing, Comcast and dozens of others. Session capture without additional authentication, with effective MFA bypass via stolen cookie.

Time to mass exploitation · ~10-14 days

CRITICAL **ScreenConnect AuthBypass · CVE-2024-1709**

Disclosed on February 19-20, 2024 by ConnectWise. Trivial authentication bypass (admin user creation via path manipulation). Mass exploitation observed in less than 24 hours — Huntress, Sophos and CISA issued urgent alerts.

Time to mass exploitation · ≤ 24 hours

CRITICAL **CrushFTP · CVE-2024-4040**

Disclosed on April 19, 2024. CrowdStrike and Rapid7 detected active exploitation within hours. Path traversal allowed reading sensitive files without authentication, including configurations and secrets. Hundreds of exposed instances compromised on the first day.

Time to mass exploitation · ≤ 24 hours

WARNING **The conclusion: anticipation is no longer a refinement**

The seven incidents share one denominator: the interval between disclosure and mass exploitation was measured in **hours or days**, never weeks. Any defense operating on a monthly cycle — or worse, quarterly — arrives afterward. The capacity to anticipate the exploit, capturing the risk signal before exploitation, becomes the **architecture most aligned with the attacker's current speed** in exposure management.

EXPOSURE WINDOW · PART 03

Anatomy of the window · from disclosure to applied patch

Four milestones define the exposure window. The interval between the second and the fourth is where the attack happens.

Operational definition. Exposure window is the interval during which a vulnerability is exploitable on an organization's surface — it begins when the exploit becomes viable and ends when the corresponding control is applied. The larger the window, the higher the probability of exploitation.

The four milestones of the window

0. **Pre-disclosure.** The vulnerability exists in the code and may be under exploitation by actors with privileged knowledge — zero-day. For the defense, this period is invisible by definition. In 2023, several incidents began in this zone.
1. **Disclosure (T0).** The CVE is published, with or without a patch available. The information becomes public — and the attacker's clock starts running at the same speed as the defense's clock.
2. **Exploit observed (T1).** The first exploit is recorded in production. In 2023, the median between T0 and T1 was **5 days** — in many cases, less than 24 hours. Here the most dangerous phase of the window opens.
3. **Patch applied (T2).** The organization applies the control on the affected asset. The Verizon DBIR observes, year after year, a **median close to 55 days** between disclosure and remediation in production.[4]

The calculation of the open window. Median exploit at **5 days**, median remediation at **55 days**. The window is the interval between the two — **~50 days per critical vulnerability**, on average, for the typical organization operating on the traditional vulnerability management cycle.

The defensive question

If the exploit arrives in 5 days and the applied patch takes 55, the defense has three technical levers to reduce the window:

- **Anticipate detection** — observe the exposure before disclosure, via threat intelligence and edge telemetry.
- **Prioritize what is being exploited** — do not treat all CVEs equally.
- **Shorten the discovery and validation cycle** — move from a monthly snapshot to a continuous flow.

THE HISTORICAL CURVE · PART 04

ASM adoption × remediation time · an inverse correlation

As ASM platforms gained adoption in mature markets, the average remediation time fell. The movement led in the US and Western Europe — and remains uneven across sectors and mid-market organizations.

The coincidence with a cause. Between 2018 and 2024, the average remediation time for critical CVEs in US and European organizations fell from about 180 to 55 days — a reduction of approximately 70%. Over the same period, the global Attack Surface Management market grew at an estimated compound annual rate of ~27-29% (CAGR), reaching a multiple of about 4.5× its 2018 size. The two movements are, to a large extent, the same movement.

Median time to remediation · critical CVEs

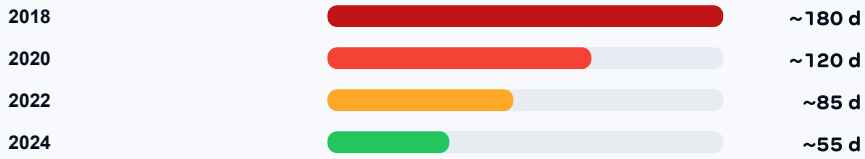

Sources: Edgescan Vulnerability Stats Report (2019-2023 series), Verizon DBIR 2024. Global median; US and European markets.

ASM market · cumulative multiplier vs. 2018

Sources: MarketsandMarkets, Allied Market Research, 360iResearch (2022-2030 series normalized to a 2018 base). Estimated CAGR: ~27-29%.

The observed correlation

The fall in remediation time was neither linear nor uniform. Markets that adopted dedicated ASM early — the US and Western Europe — concentrated most of the reduction. Industries with strong regulation (finance, defense, critical infrastructure) drove the curve. When two opposing curves cross with such consistency over six years, the correlation moves from hypothesis to a reasonable reading: the adoption of platforms specialized in continuous discovery and dynamic prioritization is the technical instrument that made it possible to shorten the window.

It is not the only factor — pipeline automation, cloud telemetry and process maturation contributed. But ASM is the component that, over those same six years, moved from a niche category to a standard budget item in corporate security programs.

Where the lag persists. Outside the early-adopter industries — finance, defense, critical infrastructure — many mid-market organizations still operate the traditional vulnerability management cycle, monthly or quarterly. Their average remediation time remains in a band the leading sectors left behind in 2020-2021, an estimated 3 to 4 years off the most advanced curve.

Adjacent ASM × specialized ASM. Many broad enterprise IT suites incorporated External Attack Surface Management modules around the core product, focused chiefly on inventory and, in some cases, discovery. A dedicated platform is organized around **detection speed** and **dynamic prioritization** — the difference that determines whether the exposure window will be days or months.

THE POINT-IN-TIME PARADIGM · PART 04

Traditional vulnerability management · a snapshot that ages

The annual and monthly scan cycle was designed for a world in which the exploit arrived in months — not days.

What the point-in-time cycle delivers — and what it misses

Traditional vulnerability management operates in cycles. The scanner is scheduled to sweep the estate at a fixed cadence — annual at many organizations, quarterly at the more mature ones, monthly at those that invest most in security. Each run is a snapshot: it records the state of the estate at the instant the scan ran.

Between one snapshot and the next, reality changes. New assets spin up in the cloud. Configurations are altered. Libraries receive new versions. New vulnerabilities are disclosed. All of this stays outside the record until the scanner's next run — and, once there, it joins the flow of findings entering the remediation queue.

This model works when the attacker's speed is comparable to the scan cadence. In 2015, with an average exploit time above 100 days, a monthly scan was reasonably aligned with the threat. In 2023, with the exploit at 5 days, the monthly snapshot shows a state that no longer exists — and the window remains open between snapshots.

The mathematical gap · open window per VM cycle

SCAN CADENCE	MAXIMUM WINDOW BETWEEN SCANS	VS. EXPLOIT AT 5 D
Annual	365 days	**73× the attacker's window**
Quarterly	90 days	**18× the attacker's window**
Monthly	30 days	**6× the attacker's window**
Weekly	7 days	1.4× the attacker's window
Continuous	< 1 day	**aligned with the threat**

A scan running once a month leaves the organization in a window six times larger than that of the median attacker. On critical assets, this ratio is the difference between detecting a vulnerability before or after exploitation.

The additional process delay. The scan is only the starting point. Triage, assignment, maintenance window, validation and re-scan extend the total cycle. In mid-sized organizations, the median between detection and remediation of a critical CVE remains around **55 days**, according to the Verizon DBIR — even when the scan operates monthly.[4]

THE COST OF THE WINDOW · PART 05

The price of an open window · in days and in dollars

The exposure window is, in essence, a financial risk metric — the larger it is, the higher the probability and cost of the incident.

241 days
average time to identify and contain a breach · IBM Cost of a Data Breach 2025[5]

$10.22M
average cost of a breach in the United States — IBM 2025

–US$0.5M
lower average cost when detection and containment occur in fewer than 200 days (US$4.93M vs. US$5.46M)

The correlation between window and cost

The IBM Cost of a Data Breach Report 2025 documents, year after year, an unequivocal correlation: **the longer the time between compromise and detection, the higher the final cost of the incident**. Organizations that detected and contained in fewer than 200 days spent, on average, US$4.93 million — against US$5.46 million at those that took more than 200 days. The difference is structural.

The reason is mathematical. An attacker with more time inside the perimeter has more assets discovered, more privileges accumulated, more data accessed and more time to prepare lateral movement and exfiltration. Each day of open window simultaneously increases the probability and the magnitude of the loss.

The exposure window, on this reading, constitutes the **central parameter** that links an organization's technical posture to its financial exposure. Reducing it is the most direct lever for reducing expected cost.

The cost of the window in the CFO's language. Each day of open window has an expected cost proportional to the probability of exploitation on that day and to the financial impact of the corresponding incident. For a critical vulnerability on a production asset, this daily cost can be estimated in thousands of dollars per day exposed — a composition that becomes the basis for calculating the ROI of remediation.

The compound effect on the portfolio

In an organization with hundreds of assets and dozens of critical vulnerabilities open at any moment, the cost of the window is cumulative. Assessing the window of a single asset is not enough — it is necessary to sum the expected cost of every open window in the portfolio.

The average reduction from 50 days to 10 days of window on a critical vulnerability corresponds, in ALE, to a **material reduction** in the expected exposure of that vulnerability. Applied to the entire portfolio, window reduction becomes the highest-return lever of the security program.

CONTINUOUS MONITORING · PART 06

The new paradigm · four principles of continuous monitoring

A different operating model, in which discovery, assessment and prioritization operate in continuous flow, replacing the scan's periodic snapshots.

Whitepaper 2026
CSURFACE

The nature of continuous monitoring. Accelerating a monthly scan to weekly keeps the same paradigm — snapshot, queue, cycle. Continuous monitoring is a **change of operating model**: discovery in flow, posture observed in real time, prioritization driven by active threat, and action triggered by event.

THE FOUR PRINCIPLES OF CONTINUOUS MONITORING

1. **Discovery in flow**
   Each new asset, each configuration change and each new component automatically enters scope the moment it appears — without waiting for the scanner's next round.
2. **Observed posture**
   The posture assessment is redone continuously, asset by asset, at each change observed in the external state.
3. **Prioritization by active threat**
   The order of the queue is determined by the real, observed threat. What is being exploited now moves to the front; the static score of a CVE published months earlier yields its position.
4. **Action triggered by event**
   Relevant changes generate alerts that feed directly into the remediation cycle — without waiting for the next periodic report.

What changes in operations

In the point-in-time model, the VM team waits for the scan, receives a report with hundreds of findings, triages by static score, assigns responsibility and opens tickets. The cycle repeats on the next run.

In the continuous model, the team works on **a flow of significant events** — new assets, posture changes, new vulnerabilities in production, evidence of active exploitation against peer organizations. Each event brings its context and priority already computed; triage moves from a sweeping exercise to a decision-making one.

Frameworks that call for continuous monitoring. The continuous model is what multiple frameworks have recommended for more than a decade. **NIST CSF** (DE.CM-8: continuous vulnerability scans), **CIS Controls v8** (Control 7: Continuous Vulnerability Management), **ISO 27001:2022** (A.8.8: technical vulnerability management on a continuous basis) and **PCI DSS 4.0** (Req. 6.3.1: identify vulnerabilities continuously) already establish the standard. The lag is in implementation, not in the standard.

THE ROLE OF CSURFACE · PART 07

How CSURFACE closes the window

Continuous discovery + threat sensor + dynamic prioritization — the three instruments through which the platform operates at the attacker's cadence.

Continuous discovery of the external surface

The platform maps the organization's external surface in flow — each new subdomain, each new certificate, each new exposure automatically enters scope the moment it becomes observable. There is no scan window: the interval between a new asset appearing and its entry in the inventory is typically hours.

It is this foundation that makes everything else possible — without continuous discovery, the threat sensor and dynamic prioritization would operate on an aging inventory.

Threat sensor · real-time threat signal

CSURFACE's threat sensor observes, in real time, the **observed exploitation activity** against the monitored surface: exploitation attempts, intelligence on published exploits, indicators of ongoing campaigns, telemetry from external sources. The result is a threat reading that reflects the present of exploitation, as it stands — beyond the static estimate of a score published months earlier.

It is this reading that decides the order of the queue: vulnerabilities under active exploitation move to the top, even when their CVSS is lower than others pending.

The contrast with the point-in-time cycle

DIMENSION	POINT-IN-TIME CYCLE	CSURFACE
Discovery cadence	Monthly / quarterly	**Continuous**
External posture	Scan snapshot	**Observed**
Prioritization	Static CVSS	**Threat sensor**
Recalibration	Next run	**In flow**
Median time-to-detect	15-30 days	**< 24 hours where test coverage exists**

The result in window terms. The combination reduces the typical exposure window of a critical vulnerability on a monitored asset from **~50 days** (traditional scenario) to the **organization's internal remediation cadence** — because detection is no longer the bottleneck. The window comes to depend only on the internal speed of applying the control.

APPLICATION AND COMPLIANCE · PART 08

The same work that protects and proves

Continuous monitoring is what multiple regulations already require. Applying it satisfies, in a single effort, both protection and regulatory evidence.

Regulatory and normative requirements

STANDARD	REQUIREMENT	ADHERENCE
FFIEC / GLBA	Operational risk management · continuous vulnerability monitoring	Met
NIST CSF	DE.CM-8 · continuous vulnerability scans	Met
CIS Controls v8	Control 7 · Continuous Vulnerability Management	Met
ISO 27001:2022	A.8.8 · technical vulnerability management on a continuous basis	Met
PCI DSS 4.0	Req. 6.3.1 · identify vulnerabilities continuously	Met
CCPA / state law	Reasonable security measures to protect personal data	Supports

The "met" reading indicates that CSURFACE's continuous monitoring delivers, in practice, the technical requirement that each standard establishes. Formal adherence depends on the contracted scope and the organization's internal policy.

Compliance as a by-product. The technical work of continuous monitoring delivers, in the same movement, the evidence that each of these standards requires — a verifiable, dated and auditable inventory of every discovery, change and remediation. The organization no longer maintains two parallel operations (one for the control, another for the evidence) and comes to operate with a single trail.

The argument for the audit committee

The audit committee asks for **evidence of a control in operation**, more than for a tool. In a point-in-time model, the evidence is the monthly scan report. In a continuous model, the evidence is the trail that shows, for each critical vulnerability, the time between the exposure appearing, the detection and the remediation.

In regulated organizations — banks, SEC-listed issuers, healthcare — this trail becomes the basis of the exposure-window KRI (Key Risk Indicator), comparable quarter over quarter.

APPENDIX · SOURCES

Methodology and sources

The external references that support the data in this whitepaper.

References

1. **CSURFACE emerging threats telemetry.** A proprietary monitoring pipeline that correlates CVE (NVD), exploitation probability index, active exploitation catalog, active-exploitation observation and campaign telemetry. Historical series 2020-2026 (partial) used in the charts on page 3 and in the executive summary KPIs. CVE funnel: NVD publication → monitored → high exploitation probability → confirmed exploitation → TTE ≤ 5 days → zero-day.
2. **Mandiant M-Trends 2024.** Mandiant's (Google) annual report on threat trends. The April 2024 edition (2023 data) records a median Time-to-Exploit of 5 days, corroborating the jump observed by CSURFACE telemetry in 2024.
3. **Google Threat Analysis Group + Mandiant.** A joint report on zero-days. It documents sustained growth of zero-days exploited in production and partial attribution to commercial surveillance vendors.
4. **Verizon DBIR 2024 and 2025.** Annual analysis of corporate incidents. DBIR 2024 reports 180% growth in breaches originating from vulnerability exploitation; DBIR 2025 projects a specific acceleration in edge devices.
5. **IBM Cost of a Data Breach Report 2025.** Annual report based on incidents investigated at more than 600 organizations across 17 countries. Global median to identify and contain a breach: 241 days.

Methodological note

The counts of fast-exploitation CVEs (TTE ≤ 5 days) cited in this whitepaper are produced by CSURFACE's emerging threats telemetry, computed over the total universe of CVEs published in the NVD each year. The methodology combines the NVD feed with active-exploitation signal, exploitation probability index and active exploitation catalog, and produces an annual funnel since 2020. The median TTE in days, published by Mandiant in the M-Trends reports, is used as a corroborating reference.

The medians of time to remediation cited follow the Verizon DBIR methodology, which computes the interval between disclosure and effective remediation in a production environment, considering the set of organizations with incidents investigated in the year.

The data specific to each case (Log4Shell, Spring4Shell, VMware vRealize, F5 BIG-IP, Citrix Bleed, ScreenConnect, CrushFTP) follow the vendors' own communications and the technical analyses published by security research teams (Cloudflare, Cisco Talos, Microsoft Threat Intelligence, Volexity).

What this whitepaper does not cover. This material is educational and descriptive. It does not constitute actuarial projection, legal or regulatory recommendation. The numbers mentioned are median values of an observed population; application to a specific organization requires modeling of its own window based on its surface, posture and telemetry.

NEXT STEPS

Close the window before the next exploit

Continuous exposure management is one of the capabilities of the CSURFACE platform — a Continuous Exposure Management platform. In an integrated 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, leaked-credential monitoring and financial risk quantification. It operates in a fully external manner — agentless and installation-free — with optional cloud, WAF and CIEM integrations available for organizations seeking greater depth of visibility.

Measure your estate's window

In a free preliminary analysis, CSURFACE maps the organization's external surface and reports the exposure window observed on critical assets.

Use the risk calculator

Estimate the financial impact of your exposure window with CSURFACE's public calculator, calibrated by market benchmark.

Read the Threat Intel whitepaper

The detailed account of dynamic prioritization — how the exploitation probability index, the active exploitation catalog and CSURFACE's threat sensor compose the real remediation queue.

Want to see the exposure window of your assets — free of charge?

Request a free analysis