surface

TELEMETRY REPORT · THREATSENSOR

The exploit clock

How much time there is between a CVE being published and the code that exploits it, measured on ThreatSensor's public telemetry. And why a score-ordered queue arrives after the fact.

CVES IN THE MONITORED BASE

928

VENDORS AFFECTED

225

MEASURED ON

17 Sep 2026

EXECUTIVE SUMMARY

What decides it, in one page

This page closes on its own. Reading only this page costs you none of the report's decisions.

Measured on 17 Sep 2026
CSURFACE · ThreatSensor

MEDIAN TIME TO FIRST PUBLIC EXPLOIT

1 day

the Cloud Security Alliance estimated five in 2023

EPSS ON THE EVE OF ENTERING KEV

0%

median across CISA's 50 most recent entries

MEDIAN DEADLINE CISA GRANTS

21 days

the market takes 43 to fix, per DBIR 2026

This report runs on a base that is both ours and public: the **928 CVEs** ThreatSensor monitors because they already carry a signal of exploitation. The base is open, asks for no sign-up, and every number here can be checked at the source.

First: the exploit lands before the plan does. The median between a CVE being published and the first public exploit in the base is one day. In 56.5% of the cases where that measure exists, the code was already there on publication day. By the conservative measure, observed exploitation, it is 21.4%, which sits in the same band as the 28.96% VulnCheck reported in January 2026.

Second: sorting by score arrives late. Across the fifty most recent CVEs to enter CISA's known exploited catalogue, the median EPSS on the eve of entry was zero. And 450 of the 917 CVEs under observed exploitation in the base score below CVSS 9, which puts them nowhere near the top of a severity-ordered queue.

Third: the regulatory deadline is faster than the practice. CISA grants 21 days at the median, and in 96 cases it granted three. DBIR 2026 measures a 43-day median to remediate, with only 26% of CISA's catalogue fully closed.

Fourth: there is no vendor to concentrate on. The 928 CVEs hit 502 products from 225 vendors. The six largest account for a third of the base, and the rest is long tail. A program that follows vendor bulletins covers the easy part of the problem.

The four converge on a conclusion about cadence rather than tooling. A seven or thirty-day cycle describes a state that changed before the report was opened. What moves the result is firing the check on the event, and reaching the remediation team with dozens of items carrying evidence instead of thousands carrying a score.

What to do with this on Monday. Pull three numbers from your own program: how many hours pass between a CVE gaining observed exploitation and your estate being re-examined for it; what fraction of last quarter's findings went through an exploitability check; and the median hours between an asset becoming exposed and the team knowing. All three fit into an afternoon and say more than any count of open findings.

The base behind this report

CVEs monitored	**928**
In CISA's KEV catalogue	**895**
With a known exploit	**261**
With a public proof of concept	**557**
Linked to a ransomware campaign	**367**
Distinct products affected	**502**

Read it with this caveat: the base is curated and leans towards KEV, which accounts for 96% of it. It does not represent the universe of published CVEs. It represents the slice that already carries a signal of exploitation, and every percentage in this report is over that slice.

PART 01

The exploit stopped waiting for disclosure

Two measures of time, because they disagree and the gap between them is itself information.

Measured on 17 Sep 2026
CSURFACE · ThreatSensor

PUBLIC EXPLOIT IN EXISTENCE

56.5%

of the CVEs with this measure already had public attack code on the day they were published

against

OBSERVED EXPLOITATION

21.4%

were already being exploited on the publication date, by the narrower criterion

Time to the first public exploit (568 CVEs)

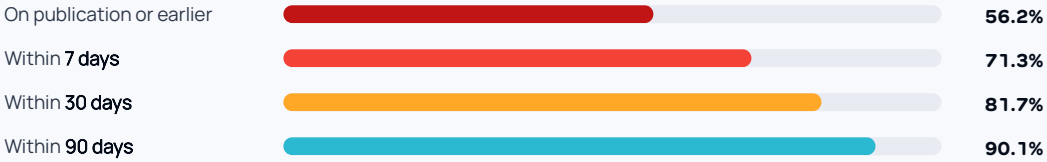

Median of minus half a day. Restricted to CVEs published in the last twelve months, the share that already had an exploit on publication rises to 62.4%.

Why the two measures disagree. The first counts the existence of public code, including a proof-of-concept repository on GitHub, a Metasploit module and an ExploitDB entry. The second counts evidence of observed exploitation, via CISA's catalogue and CIRCL sightings. Published code shows up before anyone records its use, which is why the first measure is always larger. A program planning by the second is planning by what has already been seen to happen.

The series by CVE publication year

YEAR	PUBLISHED IN NVD	EXPLOITED WITHIN 5 DAYS	EXPLOITED ON PUBLICATION
2020	21,076	43.6%	41.8%
2021	23,468	29.1%	27.2%
2022	27,554	31.6%	28.4%
2023	31,442	39.3%	28.7%
2024	39,247	32.3%	26.0%
2025	45,277	40.8%	24.2%
2026	**60,221**	**62.0%**	**52.4%**

Percentages are over the CVEs ThreatSensor monitors in each year, not over everything published. The NVD volume is the full-year count, except for 2026, which runs to the measurement date.

Discount 2026 as you read it. The year is in progress, and a recently published CVE has not had time to accumulate late exploitation. Both effects push up the share of cases that look immediately exploited. The break from previous years stays large after discounting that, but it should not be read as a settled figure.

PART 02

The score-ordered queue arrives after the attack

What EPSS and CVSS were saying about these CVEs on the eve of their entry into the known exploited catalogue.

Measured on 17 Sep 2026
CSURFACE · ThreatSensor

MEDIAN EPSS ON THE EVE OF KEV

0%

50 most recent entries; 86% below 1%

UNDER EXPLOITATION WITH CVSS BELOW 9

450

of 917 with a score assigned; 97 are medium or low

IN KEV WITH EPSS BELOW 10%

18.6%

166 of 894 CVEs in the catalogue

Backtest: which signal anticipates entry into CISA's catalogue

Temporal hold-out from 2024 onward, covering **82 CVEs** that entered KEV in the period. The question: how many of them each signal would place in its top 10% before CISA listed them.

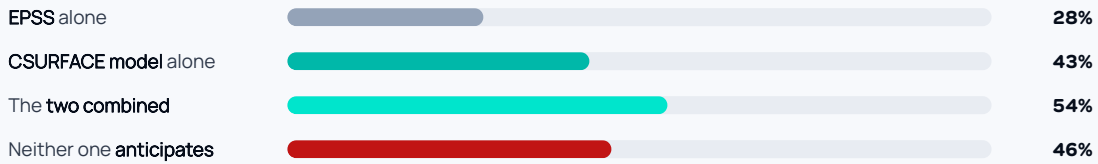

The model estimates the probability of a CVE entering KEV within 180 days and is trained without EPSS as an input, deliberately, so the two signals stay complementary. The 46% neither one catches is published alongside the rest, because a backtest that shows only its hits is useless for deciding anything.

Where exploitation concentrates

VENDOR	CVES	AVERAGE CVSS
Microsoft	**100**	8.0
Google	**65**	8.8
Apple	**44**	8.5
Cisco	**41**	8.6
Ivanti	**33**	8.5
Apache	**26**	9.3

The six largest account for 33% of the base. The remainder spreads across 502 products and 225 vendors, which is why no vendor-by-vendor inventory covers the problem.

What this means for the queue. A score is built to order things, not to discover them. While a CVE carries no public signal its EPSS is low by construction, and it sits at the bottom of any list. Once the signal appears the score climbs and the list reorders, except by then the exploitation has already started. Score-based prioritization needs a second signal, one that looks at what has not happened yet.

PART 03

The regulator's deadline is half of the practised one

What CISA grants for remediation, what the market actually takes, and the size of the gap.

Measured on 17 Sep 2026
CSURFACE · ThreatSensor

DEADLINE GRANTED BY CISA

21

days at the median between entering the catalogue and the deadline expiring, across the 895 KEV CVEs in the base

against

DEADLINE THE MARKET PRACTISES

43

days at the median to remediate, per Verizon's DBIR 2026, against 32 the year before

Distribution of the deadlines granted

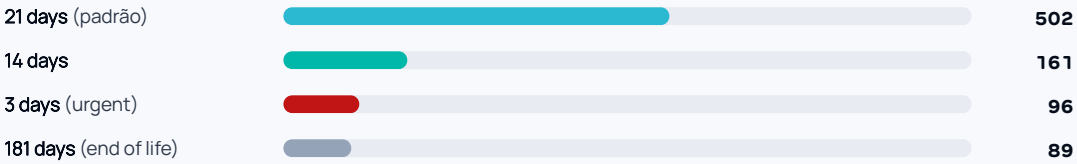

In 96 cases CISA gave three days. The shortest deadline in the base is one day. An institution working on a monthly cycle finds those cases after the deadline has already passed.

The ransomware aggravator. Of the 928 CVEs monitored, **367** have a recorded link to a ransomware campaign. This is no exotic corner of the base: it is more than a third of it, and it is the part where the interval between exploitation and financial impact is measured in hours rather than remediation cycles.

What other sources measured over the same period

Full remediation of CISA's catalogue DBIR 2026	26%
CVEs with first evidence of exploitation in 2025 VulnCheck	884
Vendors affected por essas CVEs VulnCheck	518
Catalogued by CISA in the same period VulnCheck	245
Forecast of CVEs published in 2026 FIRST	66,000
Published up to this measurement date NVD	60,221

The dispersion in our base (502 products, 225 vendors) and in VulnCheck's (672 products, 518 vendors) are independent measurements of the same order of magnitude, which is a good sign for both.

The arithmetic that remains. Between the 66,000 CVEs forecast for the year and the few hundred with observed exploitation there are two orders of magnitude. That distance is the work of prioritization, and more scanning does not cross it. Evidence that the path exists on that asset does.

csurface

METHODOLOGY AND LIMITS

How to redo these numbers, and what they do not say

Every claim in this report comes from an open base. This page exists so that it can be contested.

Measured on 17 Sep 2026
CSURFACE · ThreatSensor

Definitions that change the reading

Time to public exploit. The difference between the date of the first known public artefact and the CVE publication date. Artefact covers a proof-of-concept repository on GitHub, a Metasploit module and an ExploitDB entry. A negative value means the artefact already existed when the CVE came out.

Time to observed exploitation. The same arithmetic against the first recorded evidence of use, via CISA's catalogue and CIRCL sightings. This is the conservative measure, and the one comparable with what VulnCheck publishes.

The base is curated. It holds 928 CVEs, 895 of them in KEV. It is not a sample of the universe of published CVEs. It is the slice that already carries a signal of exploitation, and every percentage in this report is over it.

The tail is dirty and was treated. A GitHub repository's creation date sometimes predates the CVE by years, because the repository was reused. The extreme case in the base sits at minus 4,039 days. That is why this report gives medians and bands rather than averages, and why the aggregates published on the analytics page strip outliers by interquartile range.

Where each number comes from

Monitored CVE base	**ThreatSensor public API**
Funnel by publication year	**Annual funnel endpoint**
EPSS on the eve of KEV, backtest	**Analytics page**
Median remediation, CISA catalogue	**DBIR 2026, Verizon**
Observed exploitation em 2025	**State of Exploitation 2026, VulnCheck**
Historical exploit window	**Cloud Security Alliance, Apr 2026**
Volume forecast	**FIRST, Jun 2026**

Check it yourself. The base is open at csurface.io/threatsensor, with no sign-up and no query limit. The aggregates sit on the analytics page, which also explains how each chart was built. If a number in this report does not match what you measure, the error is ours and we want to hear about it.

What this report is not. It does not measure any client's estate, does not project individual risk and should not be read as a coverage promise. It is the measurement of a market phenomenon on a public base, made to support a decision about cadence rather than to describe a product.

csurface

EXTERNAL SURFACE · AUDITABLE EVIDENCE

The question stopped being how many CVEs you have. It is which ones an attacker exploits from outside.

CSURFACE audits the coverage of the scanner the institution already pays for and proves, with evidence the target itself returns, what is exploitable on the external surface. ThreatSensor, which produced the numbers in this report, is public and free, and stays that way.

External surface discovery

Maps what answers on the internet starting from the root domain, including the subsidiary's asset and whatever arrived with an acquisition. No agent, no list to supply.

Evidence the target returns

Every finding arrives with the evidence that carried it and an explicit confidence grade. When the only way to prove something would be to bring the service down, the item stays marked as probable.

Prioritization by observed risk

The queue orders itself by signals of exploitation in progress rather than theoretical severity. It is the direct application of what pages 3 and 4 of this report show.

Want this picture applied to your surface?

Get my company's preliminary analysis →