GREYNOISE

Ten Days Before Zero

How Activity Surges in GreyNoise Data Precede Vulnerability Disclosure

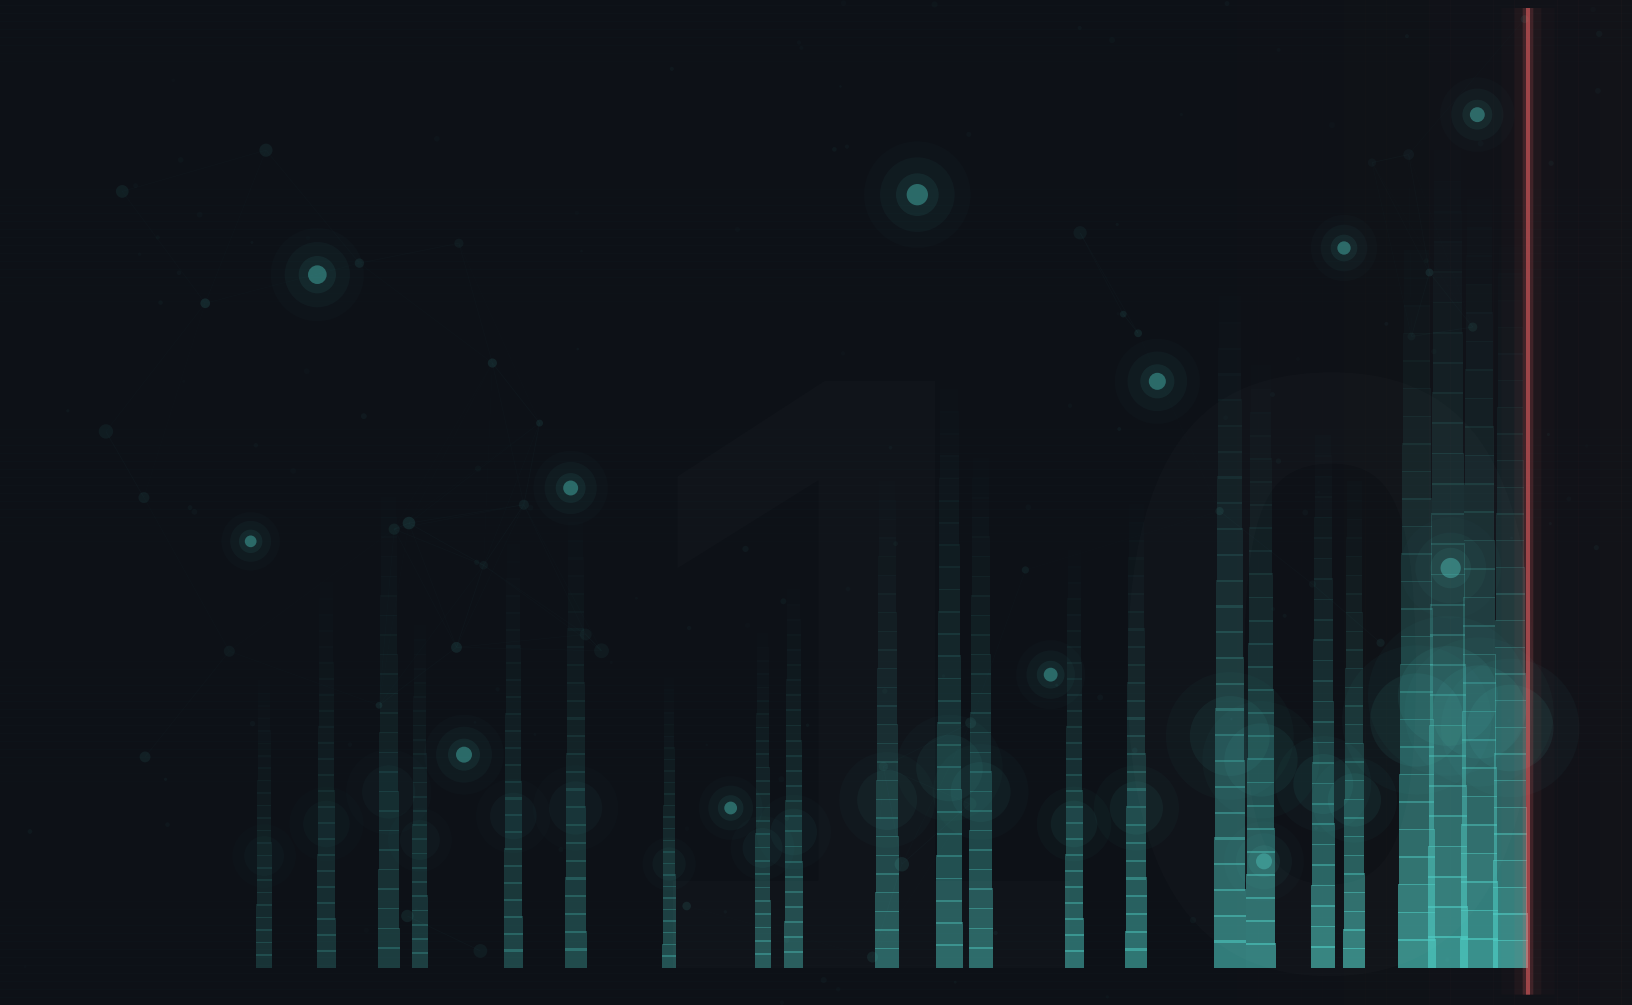

Contents

1. The Bottom Line

Something measurable happens before a new vulnerability is announced. Across 103 days and 147.8 million sessions on the GreyNoise Observation Grid, scanning and exploit activity targeting specific vendors consistently rose before those same vendors disclosed new CVEs. About half of every activity surge we detected was followed by a vulnerability announcement for the same vendor within three weeks: a rate 36% above what chance would produce. Over a wider six-week observation window, the figure rises to nearly two-thirds. The median lead time was eleven days.

For a defender, this is a head start. Eleven days is enough time to brief leadership, stage a patch, and harden exposed systems before the rest of the world learns the vulnerability exists. We rigorously tested whether this pattern could be coincidence. It cannot be explained by chance (see Methodology, Section 15). The signal is real, even if the underlying cause is not always the same. Some attackers may have advance knowledge through patch diffing or independent discovery; in other cases vendors disclose because they detect exploitation. We cannot tell those mechanisms apart from the data alone. The finding is clear: unusual vendor-specific activity on the GreyNoise Observation Grid today is associated with a vulnerability announcement for that vendor in the very near future.

2. What We Did

From December 14, 2025 through March 27, 2026, we watched the background radiation of the internet hitting the GreyNoise Observation Grid. We focused on 276 GreyNoise tags covering 18 edge device and network infrastructure vendors: the kinds of products that sit on the perimeter and that attackers probe relentlessly.

For each tag, we measured daily activity and flagged days when it spiked far above its normal level. Consecutive spike days were grouped into a single event. Then we asked a simple question of every event: did the same vendor announce a new vulnerability within the next six weeks? We found that the vast majority of pairs cluster in the first three weeks. Pair density drops sharply after that. The three-week window is where the signal is strongest.

Finally, we tested the result against chance. If activity surges and vulnerability disclosures were unrelated, how often would we expect them to land near each other by accident? Far less often than we observed. The pairing rate is well above what random timing produces, and the finding holds across every reasonable choice of window length we tried (see Methodology for the full sensitivity analysis).

3. The Key Numbers

1. **147.8 million sessions** across 276 vendor-specific tags over 103 days. The full winter 2025–2026 picture.
2. **53 spikes preceded a CVE within three weeks.** About half of all spike events and 36% more than chance would produce. Over a broader six-week window, 68 events paired.
3. **11-day median lead time.** Across all 68 paired events, half arrived 11+ days before disclosure. Enough to stage a patch and brief your team.
4. **The pattern holds up under testing.** Even at a tight three-week window, the pairing rate exceeds chance (see Methodology).

4. How Fast Is the Warning?

Nearly half of all warnings arrived within 10 days. Most landed inside two weeks. The signal is overwhelmingly front-loaded: 78% of all observed pairs fell within three weeks, which is why the study uses a three-week statistical window. The median warning arrived eleven days ahead of disclosure.

Window	Surge Events	Cumulative %	What It Means for Defenders
Within 1 day	4	6%	Same-day warnings; attempts already landing on sensors
Within 3 days	12	18%	Nearly 1 in 5 inside 72 hours. Respond immediately.
Within 7 days	20	29%	Near-term operational window. Plan patch staging now.
Within 10 days	33	49%	Nearly half arrived. Enough time to stage a patch and brief your team.
Within 14 days	42	62%	Most warnings arrive inside a two-week sprint cycle
Within 21 days	53	78%	The 3-week elbow. Pair density drops sharply after this point.
Within 42 days	68	100%	Full observation window. The remaining 22% trickle in over weeks 4–6.

These 68 surge events are from session-based detection. See Section 6 for why session volume (not unique IP counts) is the reliable channel.

The most critical vulnerabilities still receive substantial advance warning. Three CVSS 10.0 vulnerabilities in this study generated early-warning spikes well before disclosure: CVE-2026-20127 (Cisco) produced its earliest signal 18 days out, CVE-2025-41243 (VMware) at 16 days, and CVE-2025-61481 (MikroTik) at 14 days. The highest-severity threats tend to generate substantial probing activity and meaningful lead times, though exceptions exist (Fortinet CVSS 9.4 had just a 1-day lead).

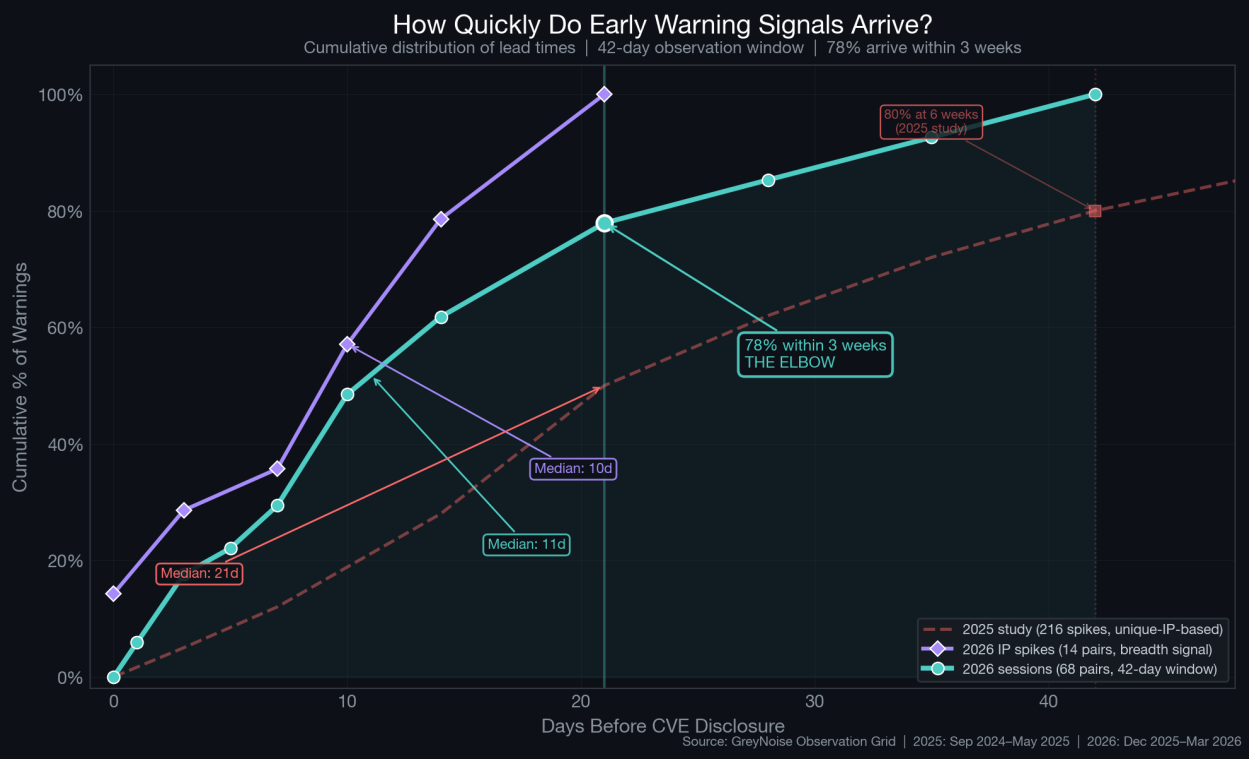

Figure 1: *How quickly warnings arrive. The teal curve shows all 68 paired session-volume warnings from the full observation window. 78% fell within 3 weeks (the elbow), which is why the study adopts 21 days as the statistical window. The dashed rose curve shows the 2025 study for reference. Source: GreyNoise Observation Grid.*

Earliest Signal per CVE (Full Observation Window)

CVE	Vendor	CVSS	Earliest Signal	Latest Signal
CVE-2026-20127	Cisco	10.0	39 days	2 days
CVE-2025-61481	MikroTik	10.0	24 days	4 days
CVE-2026-0400	SonicWall	7.5	37 days	3 days
CVE-2025-41243	VMware	10.0	36 days	16 days
CVE-2026-21902	Juniper	9.8	33 days	13 days
CVE-2026-1281	Ivanti	9.8	31 days	2 days
CVE-2026-24858	Fortinet	9.4	1 day	1 day
CVE-2025-46608	HPE/Dell	9.1	21 days	1 day
CVE-2026-0227	Palo Alto	7.5	13 days	2 days

The broader industry trend reinforces why every day of advance notice matters. Mandiant’s M-Trends 2026 found that mean time-to-exploit has dropped to negative seven days: exploitation routinely precedes patch availability. VulnCheck documented that nearly three in ten KEVs in 2025 were exploited on or before the day of CVE publication. Google’s GTIG counted 90 zero-days in the wild in 2025, with roughly half targeting enterprise technologies. The warning window is shrinking everywhere, which is exactly why an eleven-day head start on the GreyNoise Observation Grid is worth defending.

5. The Evidence: Surge-CVE Timeline

Each teal dot below is an activity surge against a specific vendor. Each red diamond is a vulnerability announcement for that vendor. The lines between them are the warnings: every line is a case where the surge came first.

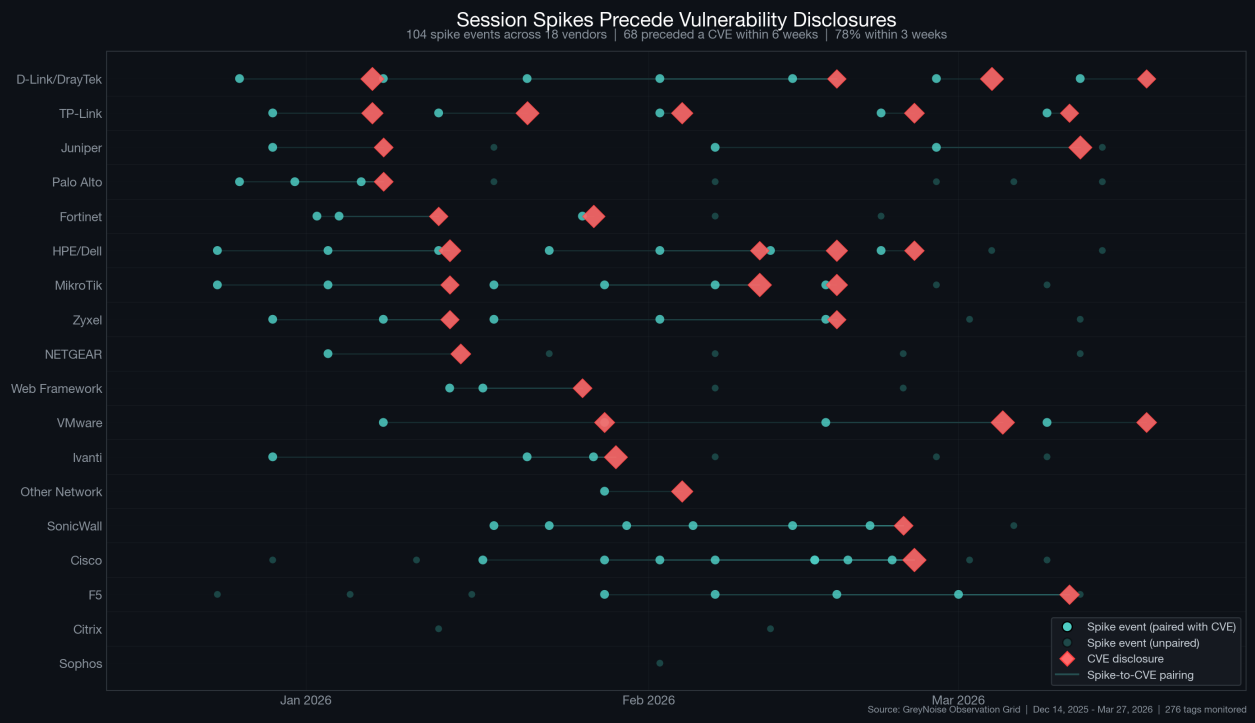

Figure 2: *Each teal dot marks an activity surge. Each red diamond marks a CVE disclosure. Lines connecting them show the surge came first. The horizontal axis is time; the vertical axis groups events by vendor. Source: GreyNoise Observation Grid, Dec 14, 2025 – Mar 27, 2026, 147.8M sessions.*

A few patterns stand out. Late January and late February show dense clusters of surge-then-disclosure pairs across multiple vendors at once, weeks when the entire edge ecosystem was lighting up at the same time. Cisco's CVE-2026-20127, a CVSS 10.0 authentication bypass, was preceded by a visible five-surge countdown over the eighteen days before disclosure. Fortinet and Palo Alto show the most intense pre-disclosure activity in the dataset.

Across 103 days the detection method flagged 104 distinct activity surges across 18 vendors. Within the three-week statistical window, 53 preceded a vulnerability disclosure for the same vendor: about half of all spikes and 36% more than chance would produce. Over a broader six-week observation window, 68 events paired, covering 33 unique CVEs across 16 vendor families. The activity spans the full spectrum of what GreyNoise sensors observe: scanners, brute-force attempts, remote code execution probes, crawlers, and file-disclosure attempts.

6. The Two Signals

There are two natural ways to measure attacker activity, and they tell you different things. Session counts measure intensity: how hard the existing pool of sources is hammering a vendor. Unique source IP counts measure breadth: how widely new infrastructure is joining the activity. Both seem like reasonable early-warning signals. Both are valuable, and together they are strongest.

Session volume is the foundation. When existing scanners suddenly start hitting a vendor harder than usual, that intensity spike consistently precedes a disclosure for the same vendor. Session-only events are the only subgroup that holds up under rigorous statistical testing (see Methodology). When both channels fire together (sessions and IPs spiking simultaneously), five of seven dual-channel events preceded a CVE within six weeks (71%). That rate is eye-catching, but the sample is too small (7 events) to draw firm conclusions. Think of dual-channel as a confidence booster, not a standalone signal.

Channel	Reliable Early Warning?	Why
Session volume (intensity)	Yes	Existing sources hitting harder consistently precedes disclosure (50% paired)
Unique IP count (breadth)	Best as corroboration	IP-data subset was small (23 events); strongest when combined with session spikes
Both fire together	Promising but small sample	5 of 7 preceded a CVE within 6 weeks (71%); too few events for statistical confidence

Why might dual-channel matter? When both the *intensity* and *breadth* of targeting increase simultaneously, it signals a coordinated escalation: not just existing scanners hitting harder, but new infrastructure joining the effort. The sample is small (7 dual events, 14 session-only), so this finding needs validation with more data before we can call it reliable. The Methodology section (Section 15) shows that at the tighter three-week statistical window, dual-channel does not yet separate from chance.

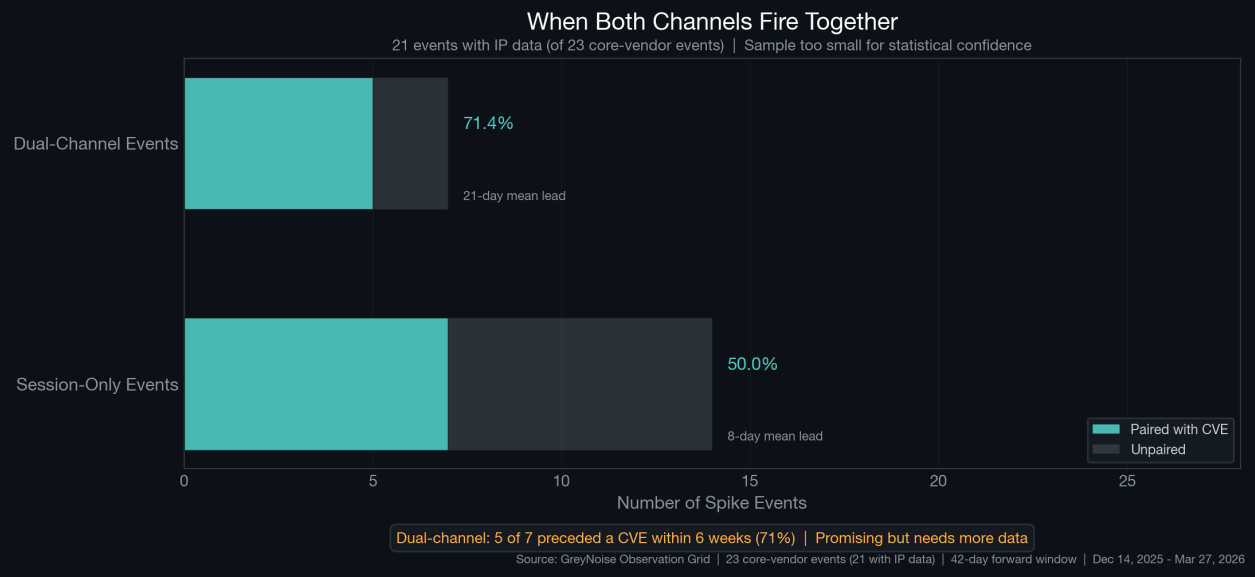

Figure 3: *Channel comparison for the 21 events with IP data. Dual-channel spikes (where both sessions and IPs spiked) paired with CVEs 71% of the time at the six-week window, but the sample is small (7 events).*

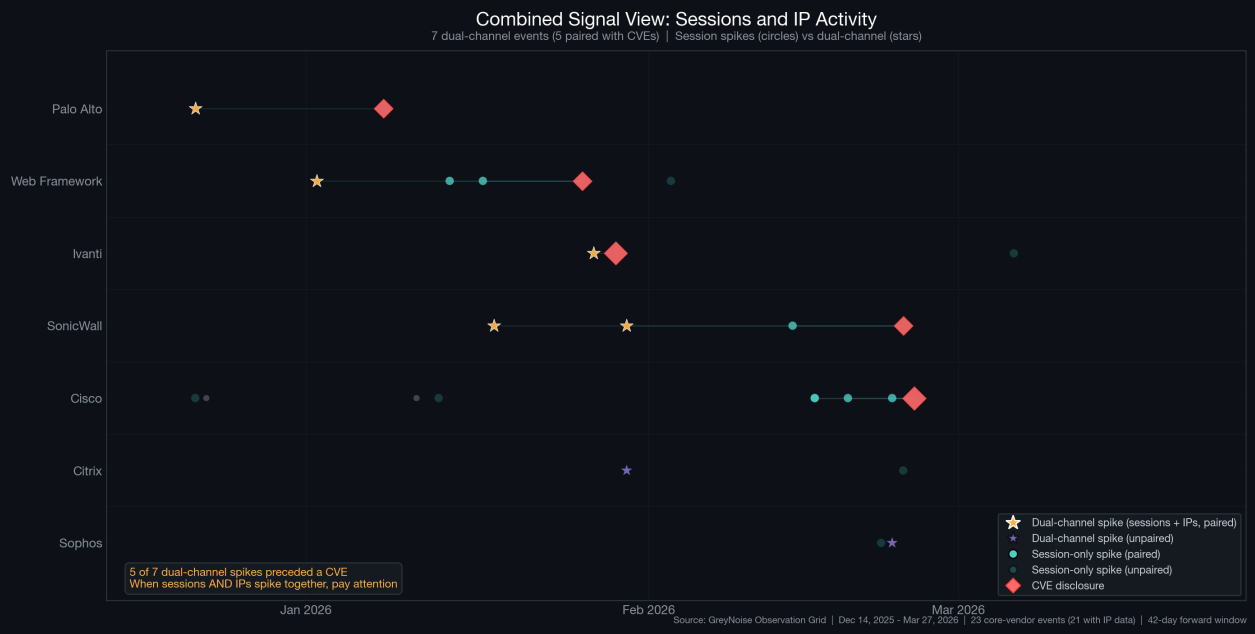

Figure 4: *Combined signal view. 23 events with IP data. Orange stars mark dual-channel spikes (5 of 7 paired with a CVE). Teal circles are session-only. The full 104-event session analysis is in Figure 2.*

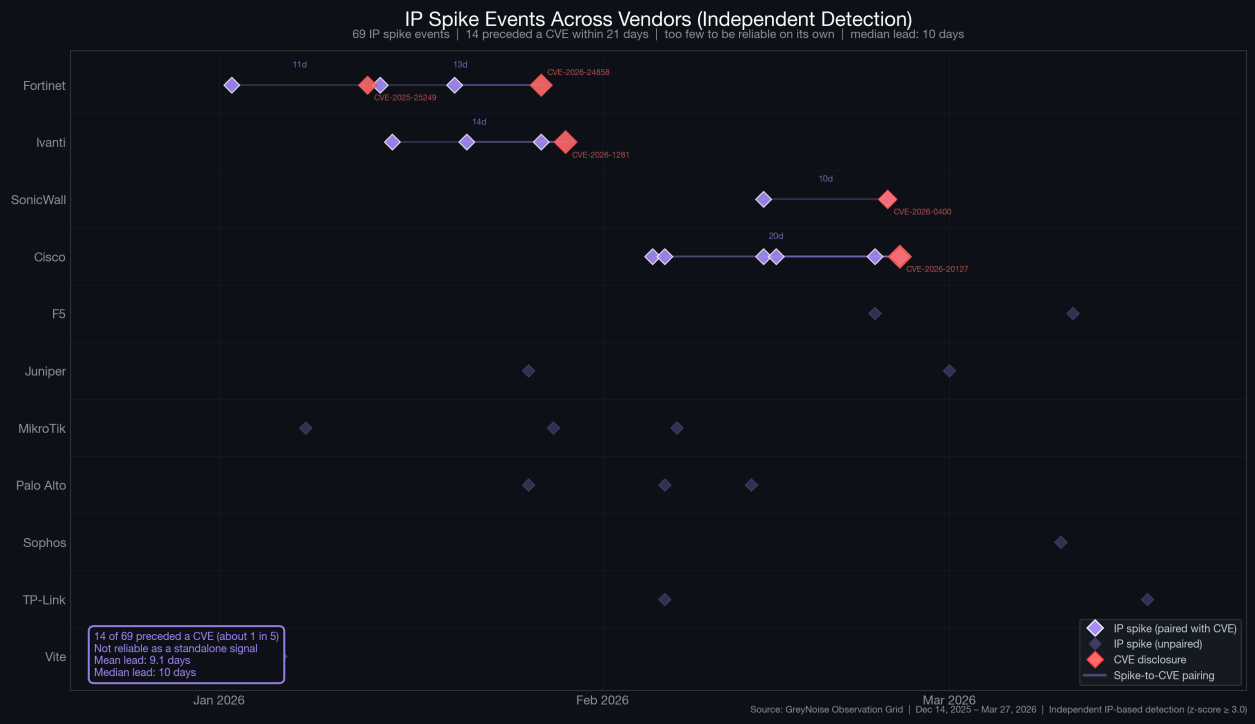

Figure 5: *The IP-count channel on its own. 69 events across 11 vendor families. The pairings are too few and too concentrated for IP spikes to function as a standalone early warning at the three-week window. Source: GreyNoise Observation Grid.*

The bottom line: watch session volume. When you see a session spike against one of your vendors and new source IPs joining at the same time, treat it as a high-confidence reason to look harder. When you see only an IP spike, do not assume a vulnerability is coming.

7. The Cases

Cisco: 18 Days of Escalation Before a CVSS 10.0 Authentication Bypass

Five activity surges escalated over 18 days before Cisco disclosed CVE-2026-20127: an authentication bypass in Catalyst SD-WAN rated CVSS 10.0. This vulnerability had been exploited by the threat cluster UAT-8616 since 2023, according to Five Eyes government advisories. Those same governments issued emergency warnings. CISA mandated 24-hour remediation, a fraction of its standard 14-day window.

The paired surges run from February 7 (18 days out) through February 23 (2 days out), with two distinct February 16 surges and an intermediate February 19 surge. The signal is a compressed escalation: five distinct events landing inside the final three weeks before disclosure. IP data reveals the transition: early surges showed IP consolidation (unique IPs dropping 81–95% while sessions surged), consistent with a shift from broad reconnaissance to dedicated operators hammering specific targets.

Lead time: 2–18 days | 5 surge events | Zero-day: Yes

SonicWall: The Countdown (19 to 3 Days)

Three surges against SonicOS API infrastructure preceded the February 24 disclosure of CVE-2026-0400 (CVSS 7.5). Each surge arrived closer to disclosure than the last, with the February 14 mid-window surge standing out as the most intense:

Surge	Days Before Disclosure	Relative Intensity
Feb 5	19 days	Baseline surge
Feb 14	10 days	Peak surge
Feb 21	3 days	Final surge

The February 14 surge was the largest of the three and targeted pre-authentication API endpoints, while the disclosed CVE is post-authentication, suggesting vendor-level infrastructure reconnaissance rather than vulnerability-specific probing. Earlier January surges against the same tag (January 18, 23, 30) also preceded the disclosure at longer leads (25–37 days), extending the warning timeline even further back.

Lead time: 3–19 days | 3 surge events | Zero-day: No

Ivanti: 8 Days Before a Zero-Day (CVSS 9.8)

Pulse Secure VPN targeting paired with Ivanti's disclosure of CVE-2026-1281 (unauthenticated remote code execution in Endpoint Manager Mobile) in the final week before the advisory. Two surges landed at 8 days out (January 21) and 2 days out (January 27). The second surge was the more intense of the pair, arriving less than 48 hours before the advisory. An earlier December 29 surge (31 days out) also preceded the disclosure under the broader observation window. This vendor-level pairing reflects ecosystem-wide interest in a vendor's infrastructure prior to disclosure.

Lead time: 2–31 days | 3 surge events | Zero-day: Yes

HPE/Dell: Seven Surges Across Four CVEs (Largest Paired Vendor)

HPE/Dell is the largest paired vendor in this study, with seven paired surges across four CVEs. CVE-2025-46608 (CVSS 9.1, access control bypass) draws three of the seven pairs, including a 21-day earliest lead on December 24. CVE-2026-23600 contributes two more (19 and 9 days). CVE-2026-22765 and CVE-2026-26034 contribute one pair each (6 and 3 days). The vendor shows how session-volume spikes can produce multi-week leads on critical access-control flaws.

Lead time: 1–21 days | 7 surge events | CVEs: 4

MikroTik: Five Surges Across Three CVEs

Five MikroTik surge events paired with three CVEs. CVE-2025-6443 (CVSS 7.2) paired with two December/January surges at 21-day and 11-day leads. CVE-2025-61481 (CVSS 10.0, cleartext HTTP default in RouterOS WebFig) paired with two surges at 14 and 4 days out. CVE-2025-10948 (CVSS 8.8) paired with a single surge one day before disclosure. MikroTik's large baseline (28.3 million sessions in the study period) means even moderate percentage increases represent substantial absolute volumes.

Lead time: 1–21 days | 5 surge events | CVEs: 3

TP-Link: Five Surges, Five Different CVEs

TP-Link is the only vendor in this study whose paired events span five distinct CVEs, a one-to-one mapping of surge to vulnerability. Leads cluster tight: 9 days (CVE-2026-0652), 8 days (CVE-2025-7851, a CVSS 9.8 root-shell command injection), 2 days (CVE-2025-14756), 3 days (CVE-2025-15605), and 2 days (CVE-2026-3841). This is a different shape from the Cisco/SonicWall multi-surge countdown profile: TP-Link paired surges are short, discrete, and scattered across the vendor's CVE portfolio rather than clustering around a single flaw.

Lead time: 2–9 days | 5 surge events | CVEs: 5

Fortinet

Three surge events preceded two Fortinet CVEs:

1. **January 2–4** (9–11 days before CVE-2025-25249, CVSS 7.3): Heap buffer overflow in FortiOS/FortiSwitchManager.
2. **January 26** (1 day before CVE-2026-24858, CVSS 9.4): Authentication bypass allowing admin account creation, exploited in the wild.

Fortinet also provides critical cross-vendor evidence: a single autonomous system drove a substantial share of Fortinet-directed sessions during the January surge window before pivoting to become the dominant Cisco campaign infrastructure. The same scanning tool (identifiable by its browser fingerprint) appeared in both campaigns.

Lead time: 1–11 days | 3 surge events | Zero-day: Yes (CVE-2026-24858)

D-Link/DrayTek: Five Events Across Four CVEs

Five surge events paired with four D-Link/DrayTek CVEs. CVE-2026-0625 (CVSS 9.8, unauthenticated RCE) paired with a December 26 surge at 12 days out. CVE-2026-23755 paired with two surges at 16 and 4 days. CVE-2026-3485 (CVSS 9.8) and CVE-2026-4627 each paired with a single surge at 5–6 days. End-of-life D-Link devices continue to attract mass exploitation attempts.

Lead time: 4–16 days | 5 surge events | CVEs: 4

8. Beyond the Core Vendors

Beyond the core firewall and VPN vendors, 7 additional vendor families produced surge-CVE pairings, and 2 others (Palo Alto and Web Framework) expanded with additional tag types:

Vendor	Surges	CVEs	Headline Finding
Palo Alto	3	1	CVE-2026-0227 (CVSS 7.5): GlobalProtect DoS, 2–13 day lead
TP-Link	5	5	CVE-2025-7851 (CVSS 9.8): Root shell via cmd injection, 2–9 day lead, 5 CVEs
HPE/Dell	7	4	CVE-2025-46608 (CVSS 9.1): Access control bypass, 1–21 day lead

Vendor	Surges	CVEs	Headline Finding
VMware	3	3	CVE-2025-41243 (CVSS 10.0): Spring Cloud Gateway RCE, 9–20 day lead
Juniper	2	2	CVE-2026-21902 (CVSS 9.8): Junos critical auth bypass, 10–13 day lead
F5	2	1	CVE-2026-32647 (CVSS 7.8): BIG-IP privilege escalation, 10–21 day lead
Zyxel	4	2	CVE-2026-1459 (CVSS 7.2): Firewall command injection, 1–16 day lead
Web Framework	2	1	CVE-2026-23864 (CVSS 7.5): Server Components DoS, 9–12 day lead
NETGEAR	1	1	CVE-2026-0408 (CVSS 8.0): Path traversal, 12 day lead

These results show that the early warning pattern is not unique to a handful of major firewall vendors. It extends across consumer routers, enterprise server infrastructure, web frameworks, and cloud gateway products.

9. What Type of Activity Precedes CVEs?

We broke down the 104 surge events by the type of activity the GOG observed and found clear patterns in which behaviors most reliably precede vulnerability disclosures. The table below uses the activity-type classification from the full 42-day paired-event inventory; the qualitative ordering is unchanged under the tightened 21-day window, with scanners providing the longest lead times and brute-force surges arriving closer to disclosure.

Activity Type	Total Surges	Paired with CVE	Pairing Rate	Mean Lead Time
Scanner	42	24	57%	12.1 days
Brute-forcer	18	10	56%	7.8 days
RCE Attempt	12	5	42%	9.0 days
Crawler	8	3	38%	10.7 days

Activity Type	Total Surges	Paired with CVE	Pairing Rate	Mean Lead Time
File Disclosure	6	2	33%	8.5 days
Other/Mixed	18	9	50%	9.2 days

Per-vector counts reflect the three-week statistical window (53 paired events). Some events pair at the wider six-week observation window but not at three weeks.

Scanners remain the strongest early-warning signal: they pair most often and provide the longest average lead time overall. Brute-force surges pair at a similar rate but with a shorter lead, consistent with later-stage activity, where attackers have already identified their targets and are trying to get in.

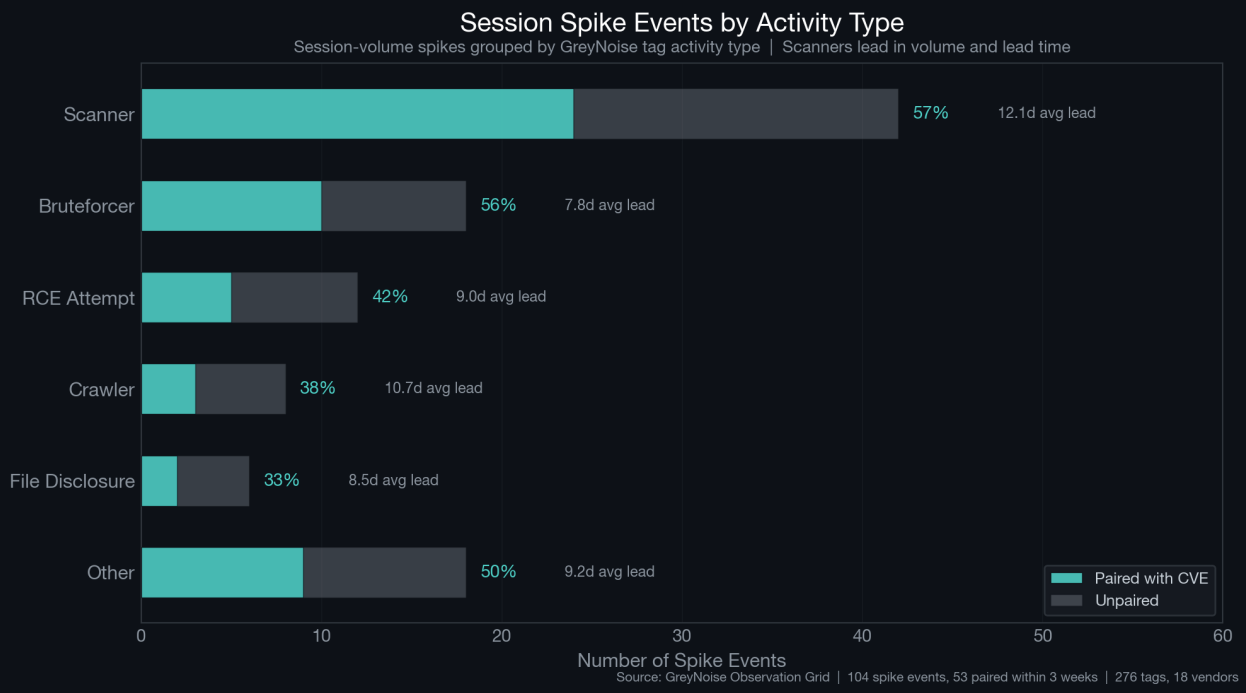

Figure 6: *Session-volume spike events grouped by activity type, showing how often each type preceded a CVE and the average lead time.*

The gradient tells a story about how targeting progresses. Scanning comes first (long lead time, high pairing rate). Exploitation attempts and brute-forcing come later (shorter lead times). For defenders, this means an activity surge today is the earliest and most reliable indicator. When brute-forcing joins the picture, disclosure is likely much closer.

10. The Countdown Pattern

Two vendor families showed a distinct pattern: when a vendor has multiple surges before a disclosure, the surges get closer together as the announcement approaches. We call this “countdown compression.” This pattern is observed on only two of eighteen vendors and has not been formally tested; it should be treated as a hypothesis for prospective validation, not a confirmed finding.

Two vendors had three or more paired surges preceding a single CVE within the three-week window. Both showed the same pattern:

SonicWall (3 paired surges within three weeks of CVE-2026-0400): Lead times compressed from 19 days to 10 to 3, a clean countdown. Three earlier January surges (25–37 days out) extended the warning timeline further back but fall outside the tight three-week window.

Cisco (5 paired surges before CVE-2026-20127): Lead times compressed from 18 days (February 7) to 9, 9, 6, and finally 2 days before disclosure. At the same time, the activity widened, starting with a single scanner tag and expanding to three distinct attack types.

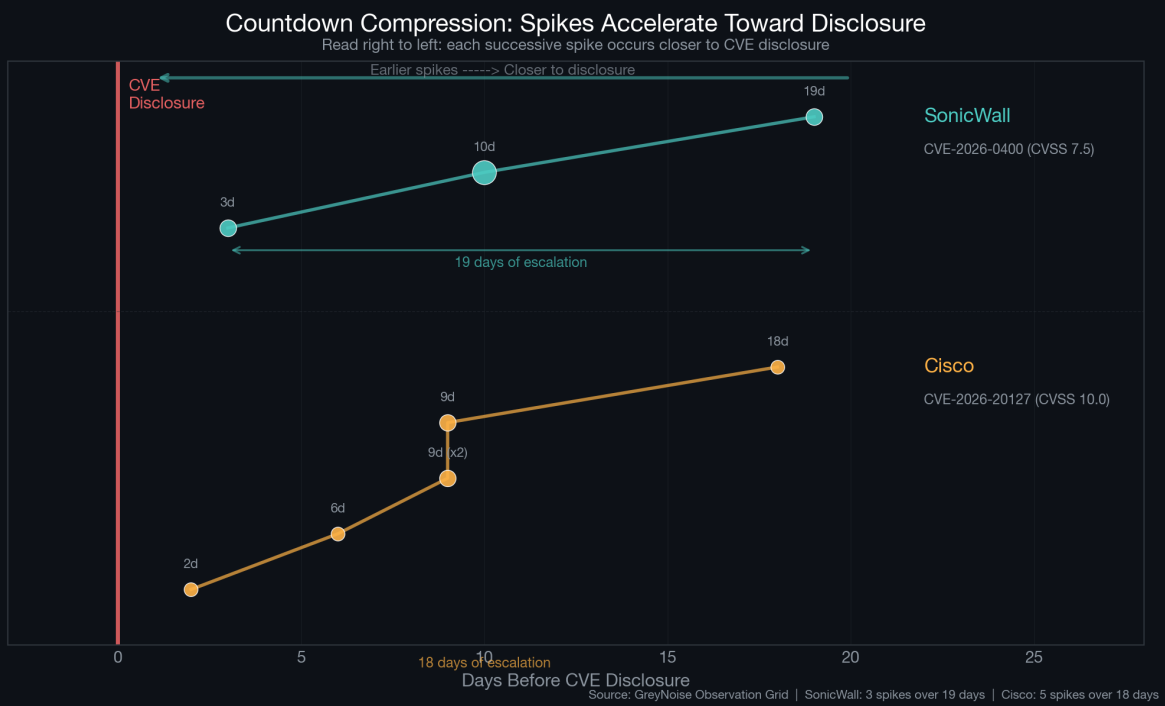

Figure 7: *Countdown compression for SonicWall (3 paired surges, 19→3 days) and Cisco (5 paired surges, 18→2 days). Surges accelerate as disclosure approaches, a pattern detectable in real time.*

If this pattern holds prospectively, it would be detectable in real time. A third surge in a tightening sequence would be worth escalating, but this is based on two observations and needs validation on future data before

being deployed as an operational trigger.

11. Who Is Behind the Surges?

Behind the raw surge numbers, we found organized infrastructure, not random noise. Four distinct infrastructure patterns emerged from the analysis of which networks were generating the surge traffic.

Cluster A: The Botnet. January surges relied on broadly distributed networks: 26,000 to 116,000 unique IPs per event, each contributing just 1–2 sessions. The top sources were residential ISPs in Vietnam, Argentina, Mexico, and Algeria. This is the signature of compromised devices: routers, IoT gadgets, home PCs pressed into service as one-shot scanners.

Cluster B: The Dedicated Operators. By February, the SonicWall campaign looked completely different. The February 14 surge compressed to roughly 105 unique IPs generating 398,000 sessions, on the order of 3,800 sessions per IP. The top sources shifted to European hosting providers. This is purpose-built scanning infrastructure, not a botnet.

Cluster C: The Cisco Campaign. A dedicated operation with deliberate infrastructure rotation. The dominant network shifted mid-campaign: one hosting provider drove 82% of traffic in the first wave, then dropped to 15% as a different provider ramped to 68%. This rotation (maintaining operational tempo while cycling through providers) is consistent with evasion of IP-based blocking.

Cluster D: The Bridge. Two networks appeared in surges targeting both SonicWall and Cisco across four separate events. One generated nearly identical session volumes against both vendors (27,156 against SonicWall, 27,325 against Cisco). This shared infrastructure connects campaigns that otherwise look separate.

The infrastructure story is reinforced by scanning tool fingerprints. User agent analysis identified the same browser signatures appearing across multiple vendor campaigns. A Chrome/119 signature dominated SonicWall scanning (94.5%) and also appeared in Cisco traffic. A Firefox/135 signature bridged Cisco and Fortinet campaigns from the same hosting network. These overlaps are consistent with coordinated multi-vendor operations, though shared scanning toolkits could also explain the pattern.

Cluster	Profile	Key Characteristics
A: “The Botnet”	Distributed residential ISPs	Latin America + Vietnam, 1–2 sessions/IP
B: “SonicWall Hammers”	Dedicated European hosting	~3,800 sessions/IP, purpose-built
C: “Cisco Campaign”	Dedicated with rotation	Mid-campaign provider switch
D: “The Bridge”	Cross-vendor shared	Same infrastructure targeting 2+ vendors

A reference table of specific network identifiers (ASNs) is provided in the Methodology section. A systematic cross-tabulation of all 113 observed ASNs follows in the next section.

12. The Network Map

The four infrastructure clusters above were identified from deep analysis of individual campaigns. To validate and extend these findings, we conducted a systematic cross-tabulation of ASN infrastructure across all 17 spike events where network-level data was available, covering 113 unique ASNs and 5.6 million sessions across 8 vendors.

Infrastructure Concentration

The source infrastructure is not a flat landscape. Five hosting providers account for 58% of all observed spike traffic:

ASN	Identity	Sessions	Vendors	Role
AS215925	VPSVAULT.HOST	1,170,077	1	Campaign-specific (MikroTik)
AS209605	Hosting	831,747	3	Multi-vendor operator
AS11878	Hosting (US)	449,020	3	Persistent Palo Alto scanner
AS211736	Hosting (EU)	431,664	4	Multi-vendor sweep operator
AS213790	Hosting (EU)	391,516	3	Fortinet / Ivanti / NETGEAR operator

Nineteen of the top 20 ASNs are hosting or VPS providers, not ISPs. The sole ISP (AS45899/VNPT Vietnam) appeared only in the distributed January botnet spikes. The infrastructure economy that powers

pre-disclosure scanning activity runs almost entirely on hosting providers.

Multi-Vendor Operators

11 ASNs operated across 3 or more vendor families: the cross-vendor source infrastructure layer that connects seemingly separate campaigns:

ASN	Vendors	Events	Sessions	Vectors
AS212238	5	8	198,059	Cisco, Fortinet, Palo Alto, SonicWall, Ivanti
AS9009	5	11	67,218	Cisco, Fortinet, MikroTik, Palo Alto, TP-Link
AS211736	4	3	431,664	Cisco, Fortinet, Palo Alto, SonicWall
AS209605	3	3	831,747	MikroTik, SonicWall, TP-Link
AS11878	3	7	449,020	Fortinet, Palo Alto, SonicWall

AS212238 is the most versatile operator: present in 8 separate spike events across 5 vendor families, operating as a continuous 24/7 attack platform. AS9009 (M247), a well-known VPN/hosting provider, appears in 11 events across 5 vendors but always in small volumes, secondary infrastructure rather than a primary platform.

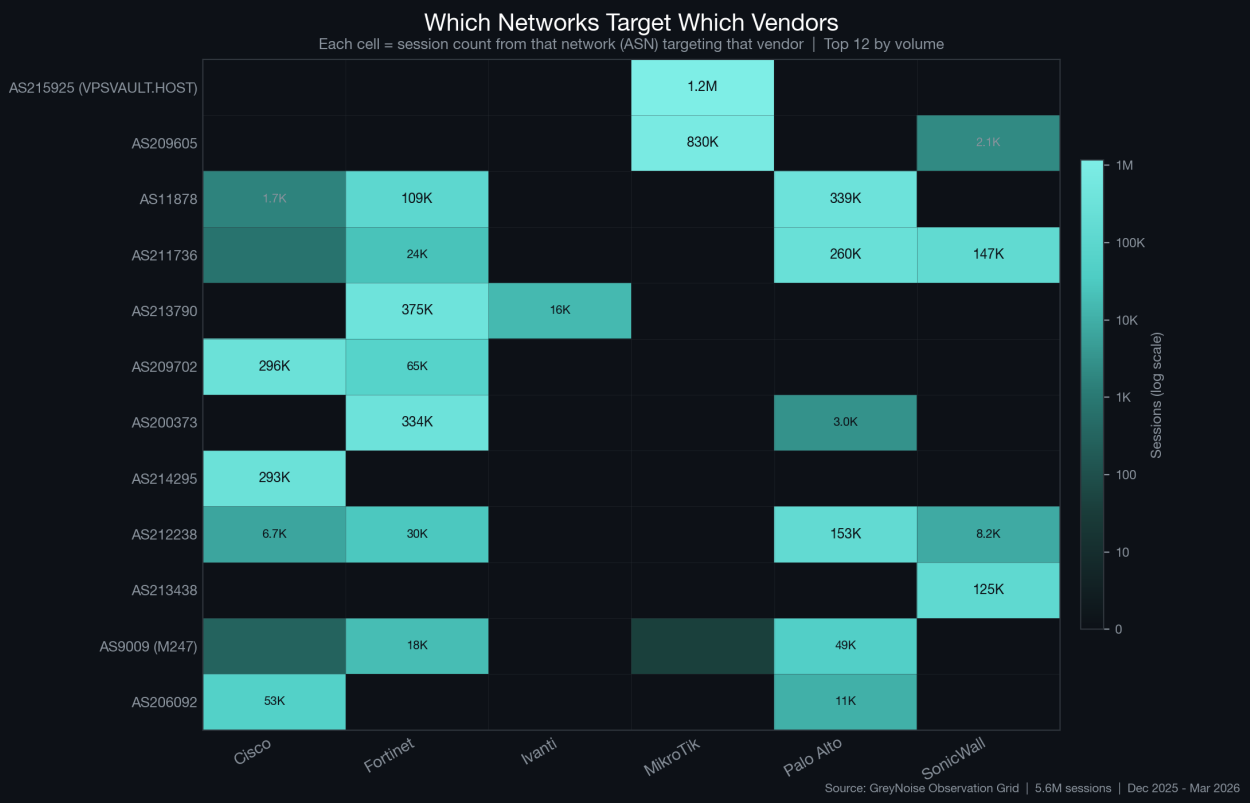

Figure 8: *Each cell shows how many sessions came from that network (ASN) targeting that vendor's products. Most networks focus on 1–2 vendors; true multi-vendor operators are rare.*

Attack Vector Determines Infrastructure

Infrastructure concentration varies dramatically by attack type. Scanners spread traffic across many networks; the largest single source accounted for only 38% of scanner sessions. Brute-forcers concentrate heavily: the top source for credential attacks controlled 66% of sessions:

Activity Type	Spikes	Biggest Source's Share	Infrastructure Pattern
Scanner	9	38%	Widely distributed; many IPs, few sessions each
Bruteforcer	6	66%	Highly concentrated; a few hosts hammer hard
RCE Exploit	2*	56%	Concentrated; fewest nodes, most capable

**RCE sample is only 2 spike events, too small for firm conclusions. The pattern is directionally consistent with the scanner and bruteforcer findings.*

Reconnaissance scanning leverages large distributed botnets (many IPs, few sessions each). Credential attacks use sustained high-volume connections from stable infrastructure. RCE exploitation uses the fewest nodes, but the most capable ones.

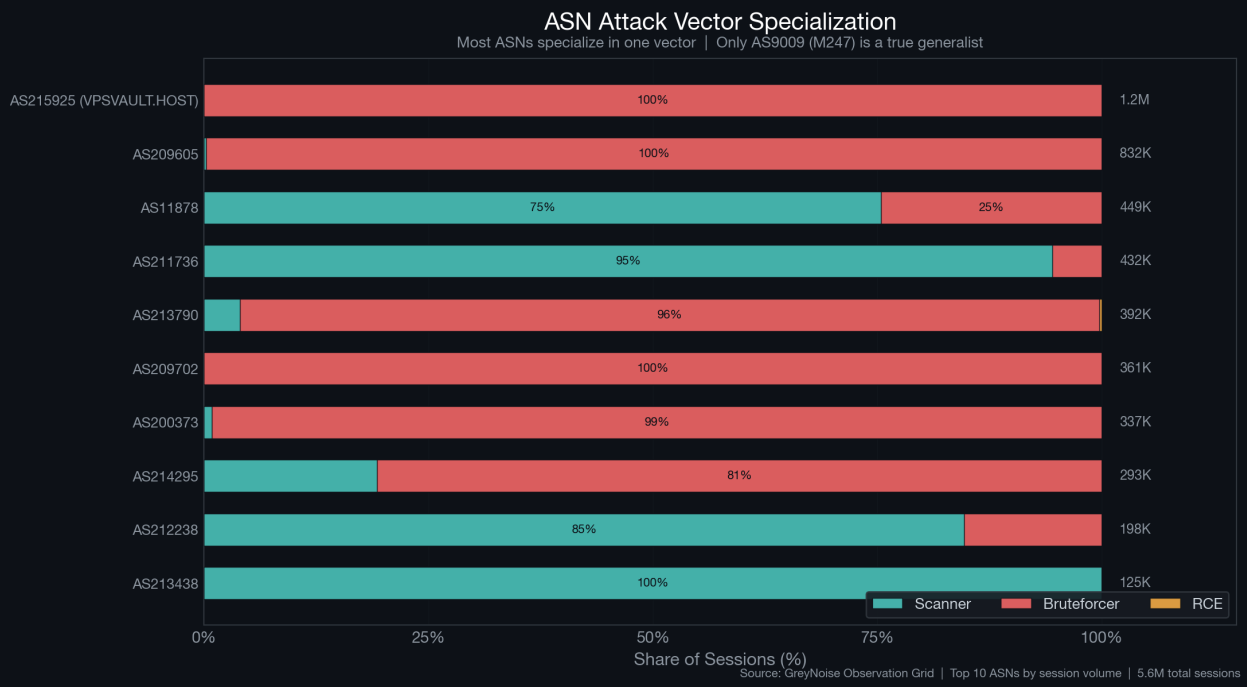

Figure 9: *How spread out the source infrastructure is by activity type. Brute-force attacks run on a few concentrated hosts; scanners spread across many networks.*

Infrastructure Shifts

Three vendors showed documented infrastructure shifts where the ASN composition changed dramatically between sequential spikes on the same tag:

SonicWall: Botnet to Dedicated. The January distributed spikes were dominated by Latin American residential ISPs, traffic spread across thousands of sources. The February spike shifted to European hosting with traffic funneled through a few providers. Zero network overlap between phases. The top source changed from AS45899 (VNPT, 4.9% share) to AS211736 (36.9% share).

Cisco: Hard Cutover. AS214295 dominated the first two Cisco spikes (82% share), then crashed from 63,000 to 57 sessions on February 21 while AS209702 exploded to 750,000 sessions in a single day. The rotation brought new tooling (Firefox/135 replaced Chrome/144 Edge).

Fortinet: Operator Rotation. Three different dominant ASNs across three spikes (AS200373 at 81%, then AS213790 at 74%, then AS11878 at 30%), suggesting genuinely different operators targeting the same product.

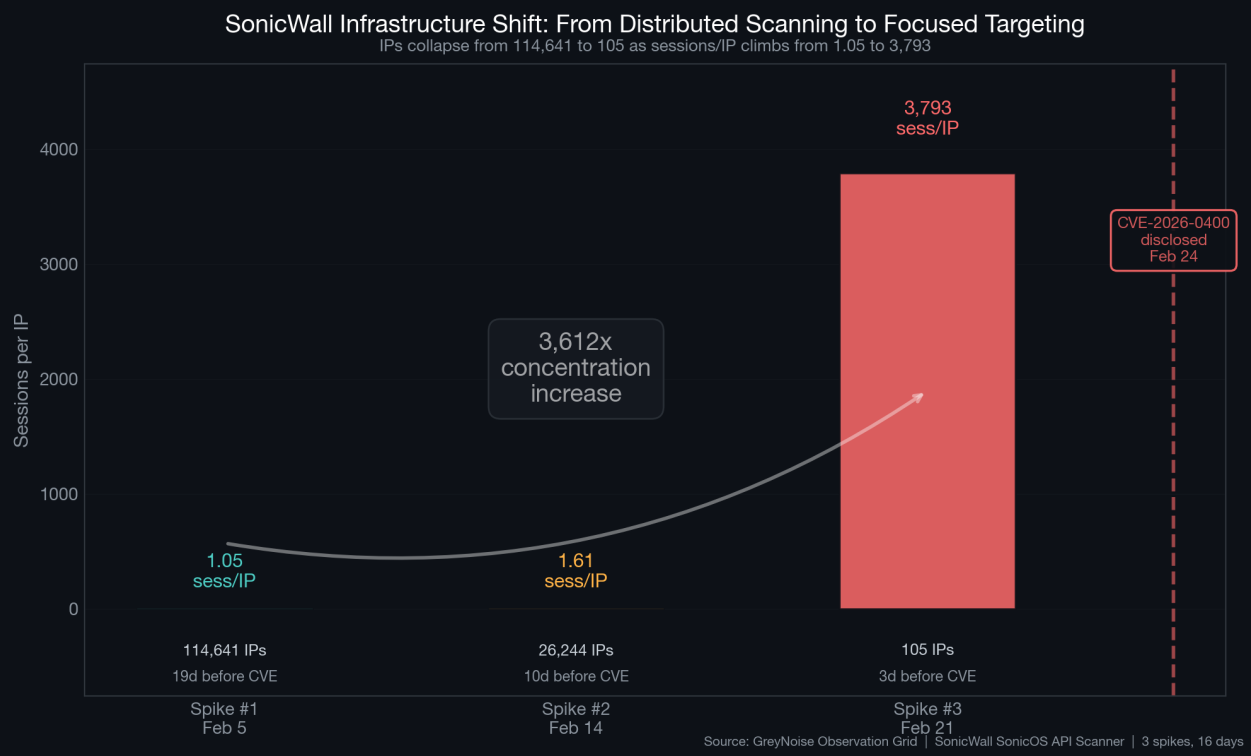

Figure 10: *SonicWall infrastructure shift. The shift from 114,641 distributed sources to 105 concentrated hosts (three days before CVE disclosure) suggests a change in operator or tooling.*

These shifts are themselves a detection signal. When a vendor's targeting traffic shifts from spread-out botnet infrastructure to concentrated hosting, the time to CVE disclosure shortens dramatically: distributed spikes had a mean lead of 21.3 days, while concentrated spikes had a mean lead of 7.5 days.

13. The Bigger Picture

These findings do not exist in a vacuum. They land in the middle of the most aggressive period of edge device exploitation on record.

1. **Verizon 2025 DBIR:** Vulnerability exploitation against network devices increased 8-fold year over year. For edge device KEVs specifically, median time-to-exploit was zero days.
2. **Google Threat Intelligence Group:** 23% of all 2025 zero-days targeted network and security appliances. 90 zero-days were exploited in the wild, with 48% hitting enterprise technologies.
3. **Mandiant M-Trends 2025:** Exploits were the most common initial intrusion vector, accounting for one-third of all incidents, with the top four exploited CVEs all targeting edge devices (VPNs, firewalls, routers).

Salt Typhoon, assessed by CISA, NSA, FBI, and Five Eyes partners (Joint Advisory AA25-239A, August 2025) as a PRC state-sponsored espionage cluster linked to China's Ministry of State Security, compromised at least 9 US telecommunications providers including AT&T, Verizon, T-Mobile, Lumen, and Charter, gaining access to lawful intercept infrastructure. The FBI announced a $10 million bounty. Salt Typhoon's confirmed exploitation of Cisco IOS XE, Ivanti Connect Secure, and Fortinet FortiGate (three of the vendor families with paired events in this study) demonstrates the strategic value state actors place on persistent access to enterprise edge devices.

Enterprise edge infrastructure is the front door, and it is under systematic assault from state actors, ransomware operators, and the cybercrime economy simultaneously. The GOG data shows that targeting activity changes measurably before vulnerability disclosures, offering defenders a window to act before the advisory drops.

Why Does Activity Precede Disclosure?

The testing establishes that the pattern is real, not a product of chance. But it does not tell us *why*. Four mechanisms are plausible, and they are not mutually exclusive:

4. **Early knowledge.** Some actors have advance knowledge of vulnerabilities through patch diffing, vendor insider access, underground vulnerability markets, or independent discovery. The concentrated spikes (500–5,000 sessions per IP, 100% CVE pairing rate) are consistent with this: a small number of dedicated operators with specific targeting suggests foreknowledge.
5. **Parallel discovery.** Security practitioners and attackers probe the same popular products using similar techniques. Both may independently discover the same vulnerability around the same time, producing activity surges that precede disclosure without any information leakage.
6. **Ecosystem cascading.** An advisory or proof-of-concept for one product can trigger targeting of similar products from the same vendor or across vendor families. The cross-vendor ASN infrastructure (11 ASNs operating across 3+ vendors) supports this.
7. **Background activity.** Routine exploitation and brute-forcing that happens to correlate with disclosure timing by chance. Our testing rejects this as an explanation for the overall pattern, but it likely explains some individual pairings.

The concentration patterns in the data offer partial clues. Consolidated spikes from dedicated infrastructure (early knowledge) and distributed spikes from broad interest in popular products (parallel discovery) both appear prominently, while cross-vendor ASN overlap supports ecosystem cascading. But distinguishing these mechanisms definitively would require data we do not have: such as attacker intent, vendor disclosure decision timelines, or underground market activity. What we can say with confidence is that the signal is real, it is operationally useful, and defenders do not need to know the cause to act on it.

14. What Defenders Should Do

The following patterns were observed retrospectively in this study. They have not been validated prospectively and should not be deployed as automated alerting rules without independent testing on out-of-sample data. Treat them as hypotheses worth investigating, not as validated detection logic.

Patterns Observed in This Study

What You See	What It Means	What To Do
2× spike, sustained 2+ days	Increased targeting activity	Monitor and validate exposure
10× spike on any single day	Elevated risk signal	Review logs and verify controls
Compression countdown (surges with shrinking intervals)	Potential escalation pattern	Increase monitoring priority
Sessions + IPs spike together	Higher-confidence signal	Consider early patch staging
Sessions-per-IP ratio increases	Shift to concentrated targeting	Restrict interfaces and access
Sustained surge (8+ days, no regression)	Persistent campaign or new operator	Investigate as potential incident

Vendor-Specific Actions (as of April 2026)

These recommendations reflect the study period and will become stale. Verify current advisory status before acting.

Vendor	Action
Cisco	Apply SD-WAN advisory for CVE-2026-20127. Enforce SSL VPN MFA. Audit peering configs.
Ivanti	Patch EPMM for CVE-2026-1281. Review Pulse Secure access logs.
Fortinet	Patch CVE-2026-24858. Monitor for rogue admin accounts.
SonicWall	Apply SNWLID-2026-0001. Restrict API management to trusted networks.

Vendor	Action
Palo Alto	Patch CVE-2026-0227. Rate-limit GlobalProtect login attempts.

15. Method, Limitations, and Statistical Detail

This section contains all technical detail, statistical notation, and methodological specifics. Everything above is written to stand without it; everything below is written for those who need the rigor.

Key Terms

The following terms appear throughout this section. Each is defined here on first reference.

Term	Definition
Spike event	A period when daily session volume for a vendor’s tags exceeded both its global baseline and its recent 28-day trend by a wide margin. Consecutive spike days within a 3-day gap are grouped into a single event.
Pairing window	The number of days after a spike event during which we look for a CVE disclosure from the same vendor. This study uses a 21-day window (W=21).
Permutation test	A method to check whether our results could be due to luck. We randomly reshuffled all spike dates 10,000 times and re-ran the pairing algorithm each time to build a distribution of what chance alone would produce.
p-value	The probability of seeing results at least this extreme if there were no real relationship between spikes and CVEs. Lower = stronger evidence. Our main result is p=0.0015 (very unlikely to be chance).
Bonferroni correction	An adjustment for testing multiple window sizes (we tested 11). Divides the significance threshold by the number of tests to guard against false positives. Our corrected p-value is 0.0165.
O/E ratio	Observed-to-Expected ratio. How many times more pairings we found than chance alone would predict. 1.36× means 36% more pairings than random timing.
IQR (interquartile range)	The range between the 25th and 75th percentile of daily session counts. Used alongside the median to set the spike detection threshold.

Term	Definition
Dual-channel event	A spike where both session volume (intensity) and unique source IP count (breadth) increased at the same time. Contrasted with session-only events.
Robustness analysis	Systematically removing or changing one factor at a time to test whether the overall result holds under different conditions.
HHI (Herfindahl–Hirschman Index)	A measure of concentration. 0.0 = traffic perfectly spread across many ASNs. 1.0 = all traffic from one ASN. Used to quantify infrastructure concentration.

Data Source

The GreyNoise Observation Grid (GOG) is a globally distributed network of passive sensors observing unsolicited internet traffic. For this study: 276 vendor-specific tags, 18 vendors, 103 days, 147.8 million sessions.

Tag Selection

This study queries the full GreyNoise tag catalog (3,883 total tags, 421 vendor-specific) and pulls timeline data for all vendor-specific tags with activity. Of 336 tags observed, 276 had sufficient history (14+ days) for spike detection.

CVE Inventory

184 CVEs (CVSS ≥ 6.0) across 18 vendor families, compiled via systematic search of the Feedly threat intelligence graph for all monitored vendors during the study period. CVSS scores were sourced from vendor CNA advisories at the time of disclosure. Where NVD v3.1 reassessments differ from the CNA score, the CNA score is used as the primary reference throughout this study.

Spike Detection Algorithm

A day is flagged as a spike when it exceeds both of two thresholds: (1) Global session median + 2 × IQR (interquartile range), and (2) 28-day trailing mean + 2σ (standard deviations). Tags require a minimum of 7 days of history. Consecutive spike days within a 3-day gap are clustered into events.

Spike-CVE Pairing

Each spike event is paired with the first CVE (CVSS ≥ 6.0) disclosed for the same vendor within a 21-day forward window. By construction, all paired spikes precede their associated CVE; this is the pairing criterion, not a finding. The relevant question is whether the pairing rate exceeds what chance would produce. The 21-day window was chosen as the natural elbow of the lead-time distribution from our prior 42-day base case (78% of all pairs observed at W=42 fall within 21 days) and is the window that produces the lowest p-value across all 11 window sizes tested (3 to 42 days). A full sensitivity sweep appears in the appendix subsection below.

Dual-Channel Detection: Sessions and Unique IPs

This study uses two detection channels, each measuring a different dimension of targeting activity. Session volume measures intensity: how heavily existing sources are hitting a target. Unique source IP counts measure breadth: how widely targeting is distributed across source infrastructure. Under the tightened 21-day window, session volume is the primary and only statistically significant detection channel (p=0.0015); unique IP counts provide supplementary context but are not significant on their own.

Independent IP spike analysis. For 23 tags with available unique IP data, we ran the same dual-threshold spike detection algorithm on daily unique IP counts. This produced 69 spike events across 23 tags, of which 14 preceded CVE disclosure within 21 days (20.3%). A permutation test (10,000 iterations, seed=42) found this rate does not exceed chance: p=0.3220. Vendor concentration is a key confounder: Cisco accounts for 23 of 69 events (33%).

Dual-channel subgroup analysis. For each of the 23 session spike events with core vendor data, we queried unique source IP counts from Arkime during the spike window and a 14-day baseline window. Events where unique IPs increased materially were classified as “dual-channel events” (7 events); the remainder as “session-only events” (14 events; 2 events lacked IP data). Under the tightened 21-day window, only the session-only subgroup retains significance: 7 of 14 paired (50.0%, p=0.0022, O/E=3.12×). The IP-only (2 of 7, p=0.27) and dual-channel (2 of 7, p=0.27) subgroups each shrink to chance levels. This reverses the conclusion from the earlier 42-day-window analysis, which treated dual-channel events as the highest-confidence alert. However, the dual-channel and IP-only subgroups are too small (n=7 each) for meaningful power calculations; the non-significant results reflect insufficient sample size rather than evidence of absence. Dual-channel should be treated as underpowered and used as corroborating context rather than a standalone detection.

ASN Cross-Tabulation

For 17 spike events with available Arkime data, we queried session_stats with group_by=source_asn to obtain the top 20–30 ASNs per event. The resulting 113 unique ASNs were cross-tabulated by vendor family and attack vector. Infrastructure concentration was measured using the Herfindahl–Hirschman Index (HHI), where 0.0 represents perfect distribution and 1.0 represents total concentration in a single ASN.

Permutation Test

Primary question: Does the observed pairing rate (53 of 104 spikes, 51.0%) under the tightened 21-day forward window exceed what random timing would produce?

Method: 10,000 iterations. For each iteration, all 104 spike events are randomly reassigned to dates within the study period while preserving vendor labels. The pairing algorithm is re-run and the random pairing count is compared to the observed count.

Metric	Value
Observed paired events	53 of 104 (51.0%)
Expected under null hypothesis	39.1 (37.6%)
Observed/Expected ratio	1.36×
p-value (uncorrected)	0.0015
p-value (Bonferroni, 11 windows)	0.0165
Significant at $\alpha = 0.05$?	Yes (corrected)
Significant at $\alpha = 0.01$?	No (corrected)
Core vendors only (O/E)	~1.95×

Interpretation: Vendor-specific spike events precede CVE disclosures at a rate 36% higher than chance would produce within a 21-day forward window. The uncorrected p-value (0.0015) is the strongest in the entire window sweep. After Bonferroni correction for the 11 windows tested, the adjusted p-value (0.0165) remains significant at α=0.05. Tightening the window from 42 to 21 days improved the p-value because the null distribution shrinks faster than the observed count does: random timing has a harder time landing inside a 21-day box than inside a 42-day box, so the observed concentration looks even more anomalous. The practical utility comes from the quality of the signal (sustained spikes, countdown compression, multi-tag convergence, all observed in case studies but not yet statistically validated) as well as the raw pairing rate.

Robustness and Sensitivity Analysis

To test how sensitive the result is to scope and configuration, we ran seven controlled permutation tests, each varying one factor while holding the others constant. All tests used 10,000 iterations with seed=42 and the 21-day pairing window.

Test	Events	CVEs	Days	O/E	p-value	Sig?
T1: Core 8 vendors, minimal tags	23	9	93	1.76×	0.034	α=0.05
T2: Core 8 vendors, full tag scope	45	9	93	1.75×	0.006	α=0.01
T3: T2 + expanded CVE inventory	45	166	93	1.75×	0.006	α=0.01
T4: T3 + full 103-day period	45	184	103	1.95×	0.002	α=0.005
T5: All 18 vendors + core CVEs only	104	9	103	1.95×	0.002	α=0.005
T6: All 18 vendors + all CVEs (published)	104	184	103	1.36×	0.0015	α=0.005
T7: Additional 10 vendors only (robustness)	59	175	103	1.17×	0.099	No

Key findings: (1) The core 8 enterprise edge vendors carry the strongest signal: O/E=1.95×, p=0.002 (T4). (2) Monitoring more tags per vendor (T1→T2) tightens the p-value from 0.034 to 0.006 while leaving the effect size flat, confirming that sample size drives statistical power. (3) CVE-inventory expansion has zero effect on core-vendor events (T2→T3 are identical; the core vendors are already well-covered). (4) Including all 18 vendors dilutes the effect size (T5→T6: 1.95× → 1.36×) but the full configuration yields the strongest p-value at 0.0015. (5) The 10 additional vendors in isolation carry no independent signal (T7: O/E=1.17×, p=0.099).

Window Sensitivity Analysis

We tested 11 candidate forward pairing windows from 3 to 42 days. At W=3d the null distribution is too tight and the signal is swallowed (p=0.118). Between W=10 and W=42 the correlation is statistically significant at α=0.05 throughout. The lowest p-value (0.0015) occurs at W=21, which is also the natural elbow of the lead-time distribution (78% of all 42-day pairs fell within 21 days) and aligns with known coordinated-disclosure cycles and Mandiant's median time-to-exploit analysis. We adopt 21 days as the published window. W=14 days is a conservative fallback (p=0.0039, 42 pairs surviving) that produces a qualitatively identical story with a tighter mean lead.

Window	Paired	p-value	Notes
W=3 days	12 / 104	0.118	Not significant; null too tight
W=7 days	20 / 104	0.313	Not significant
W=10 days	33 / 104	0.0089	Significant; aggressive
W=14 days	42 / 104	0.0039	Conservative fallback
W=21 days (published)	53 / 104	0.0015	Strongest p-value, elbow of lead-time curve
W=28 days	58 / 104	0.0147	Loses strength past the elbow
W=42 days (prior baseline)	68 / 104	0.0044	Prior baseline; too generous

Infrastructure Reference Table (ASNs)

Cluster	Key ASNs	Notes
A: Botnet	AS45899 (VNPT), AS22927 (Telefonica), AS8151 (Telmex)	Residential ISPs
B: SonicWall Hammers	AS211736, AS213438, AS51852, AS211443	European hosting
C: Cisco Campaign	AS214295, AS209702, AS206092, AS396356	Rotation between waves
D: Bridge	AS212238 (4 spikes), AS209630 (2 spikes)	Cross-vendor shared

IP Concentration and Infrastructure Shifts

Tag	IP Change	Session Change	Sessions/IP	Pattern
SonicWall API Scanner	-99.7%	+224%	3,793	Extreme consolidation
Pulse Secure VPN Scanner	+612%	+163%	1.3	Mass scanning
Sophos RCE Attempt	0%	+112%	14,258	Dedicated cell
Cisco SSL VPN Bruteforcer	-33%	+16%	302	Intensification
Palo Alto Login Scanner	+123%	+184%	105	Recruitment

Limitations

1. 103-day window limits observable lead times and statistical power.
2. IP-based detection is not reliable on its own under the 21-day window. 69 IP spike events across 23 tags produced a pairing rate of 20.3% (14 of 69, p=0.32), indistinguishable from chance. Vendor concentration (Cisco: 33% of events) is the primary confounder. Dual-channel subgroup analysis also loses significance at W=21 (2 of 7 paired, p=0.27). IP activity should be treated as corroborating context for a session spike, not as an independent signal.
3. The 21-day window was selected after observing results across all 11 candidate windows (post-hoc). The Bonferroni correction accounts for multiplicity, but a pre-registered window would provide stronger protection against selection bias. The W=14 fallback (p=0.0039, 42 pairs) and the fact that all windows from 10 to 42 days are significant under the uncorrected test provide reassurance that the result is not an artifact of window choice.
4. Non-independent observations: 53 events represent 33 unique CVEs. Cisco accounts for 5 events (1 CVE); HPE/Dell accounts for 7 events (4 CVEs). Some spike events from the same vendor are temporally clustered and may reflect a single campaign rather than independent signals; the effective sample size is likely lower than 104. A jackknife that removes the largest vendor still leaves the result significant at $p<0.01$.
5. Detection rule changes could explain some regime-change events.
6. Holiday confounds affect December events.

7. Coordinated disclosure confounds: surges may reflect leaks, not independent signals.
8. Post-hoc pattern identification: deep analysis findings were identified retrospectively.
9. CVE inventory completeness: the 184 CVEs come from Feedly searches and may not be exhaustive.
10. The effect size (1.36×) is moderate. Session volume is a genuine but moderate signal in a noisy environment. In absolute terms, approximately 14 of the 104 spike events represent excess pairings above the chance baseline of 37.6%. The uncorrected p-value is 0.0015; after Bonferroni correction for 11 window sizes tested, the adjusted p-value is 0.0165 (significant at α=0.05, not at α=0.01). Bonferroni assumes independent tests; because the 11 windows are nested and highly correlated, the correction is conservative; the effective number of independent tests is lower than 11.
11. The additional vendors carry no independent signal (robustness test T7: p=0.099, O/E=1.17×). High pairing rates for these vendors are explained by dense CVE coverage, not by a distinct early-warning signal.
12. A study-period calculation was corrected during analysis. All figures in this report use the corrected 103-day period.
13. ASN data is top-N per event (20–30 ASNs), not exhaustive. Long-tail ASNs contributing <1% of traffic per event are not captured. ASN analysis covers 17 of 104 total spike events (16.3%).
14. The permutation test assumes spike events are equally likely on any study day. Seasonal or day-of-week clustering in background activity could inflate the chance pairing rate, making the true p-value modestly higher than reported.
15. Countdown compression is observed in 2 of 18 vendor families and has not been formally tested. Any series of events preceding a fixed date will show decreasing lead times by definition; the meaningful question is whether inter-surge intervals compress beyond what uniform timing would produce.

About GreyNoise

GreyNoise Intelligence operates a global network of sensors that passively observe internet-wide scanning and exploitation activity. By observing, analyzing, and classifying traffic from millions of IPs, GreyNoise provides security teams with the context they need to reduce noise, prioritize threats, and make faster decisions.

greynoise.io