

FortiBleed Is Real — and It's Not About Fortinet

A cross-vendor measurement of edge credential exposure. The credentials behind this week's headlines reflect an industry-wide pattern in remote-access credential hygiene — common to every major platform, not specific to the one in the news. We measured it the same way across five major VPN platforms.

Beacon Technology Group · HUMINT Team · June 19, 2026

SUMMARY

This week two firms reported a credential-harvesting campaign — branded “FortiBleed” — said to affect tens of thousands of internet-facing Fortinet firewalls and VPNs. The core of it is real, and worth stating plainly: the credentials are genuine, many still work, and the activity is ongoing. Whether a new vulnerability is involved is disputed among the researchers tracking it — but either way the result is the same: exposed, working credentials, of a kind not unique to any one vendor. On June 18 the U.S. Cybersecurity and Infrastructure Security Agency issued an advisory on the campaign; its recommended response — reset credentials, enforce phishing-resistant MFA, lock down management interfaces — is worth acting on without delay, and is, tellingly, all credential hygiene and configuration rather than any product fix.

Update — June 19, 2026. After publication, researcher Kevin Beaumont documented logged configuration exports from affected devices and disputed the characterization of the data as merely prior-breach material. The original access vector is now actively contested and may involve an as-yet-unconfirmed flaw. None of this changes the measurement below: cross-vendor credential exposure is footprint-driven, and Fortinet sits fourth of five regardless of how the configurations were obtained. We flag the open question rather than resolve it.

That last point is the story. If FortiBleed required no Fortinet-specific vulnerability, then the exposure it describes should not be Fortinet-specific either — it should appear, at comparable scale, across every vendor whose appliances sit on the public internet and authenticate with reusable credentials. We tested that directly against our own credential-exposure corpus, measuring the number of distinct organizations with remote-access portal credentials circulating in commodity infostealer data, using one fixed method applied identically to five major VPN platforms.

The result: Fortinet is not an outlier. Across five platforms the per-vendor counts range from about 1,700 to 3,600 distinct organizations, and Fortinet — the vendor in the headlines — sits fourth of five: three of the platforms we measured show higher exposure, one shows less. The variation tracks how much remote-access surface each platform has deployed over the years, not how secure it is. Separately, of the marquee “victims” named in the coverage, the four we sampled did not run the attack surface the campaign exploits at all. The dominant signal in the data is not vendor-specific weakness — it is the breadth of credential exposure across every major remote-access platform.

THE CLAIM, AS PUBLISHED

What FortiBleed is reported to be

Before measuring anything, the record — stated in the reporting parties’ own figures and dates, with no characterization added.

The campaign was first surfaced over the weekend by independent researcher Bob Diachenko, who found an exposed server holding working Fortinet credentials. It was then analyzed and published this week by Hudson Rock and SOCRadar, and the dataset was independently confirmed as legitimate by researcher Kevin Beaumont. As reported by *TechCrunch* on June 17, the operation uses automated scanning plus previously known passwords to reach exposed devices; it abuses no unknown vulnerability.

Fortinet’s position, provided to *TechCrunch*, is consistent with that: the company characterized the data as a reshare of credentials from prior incidents combined with brute-forcing, and stated it is **“not related to any recent incident or advisory.”** On the central technical question — whether this represents a new vulnerability in

Fortinet’s products — there is no settled answer: Fortinet says it does not, while other researchers point to evidence of an as-yet-unexplained access vector. That question is unresolved, and it is separate from what this analysis measures.

Credential stuffing Replaying username/password pairs already known from prior leaks or infostealer logs against a login portal, to find which still work. Requires no software vulnerability — only that a valid credential was never rotated.

Config exfiltration Obtaining a device’s configuration file, which can contain stored credential hashes and service-account binds. Distinct from credential stuffing, and the part of FortiBleed whose original access vector remains officially unconfirmed.

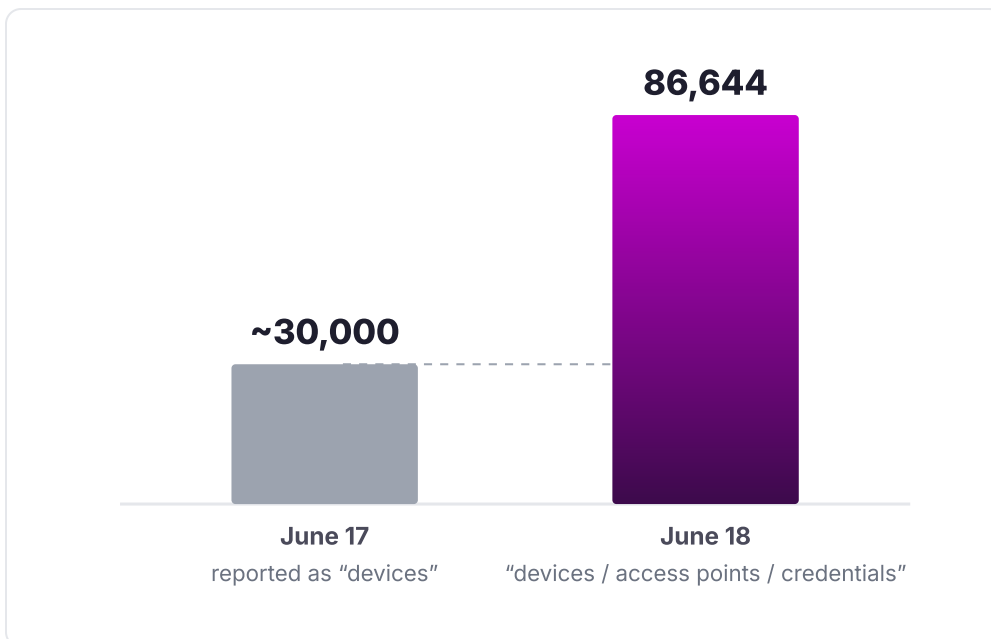
These mechanisms have different technical causes but converge on the same outcome: valid credentials that authenticate to exposed remote-access infrastructure. That shared endpoint — a working credential, not a broken device — is what this analysis measures.

The published figures — dated

The reported scale moved within the news cycle, and the unit being counted moved with it. We record both as published:

EXHIBIT A One campaign, two reported magnitudes, two days apart

SOCRadar’s reported figure, as carried by *TechCrunch* on June 17, against the figure on SOCRadar’s own page on June 18. Values and unit labels are as published by the source.



A measured quantity is stable under its own definition. SOCRadar's figure was reported as more than 30,000 on June 17 and stands at 86,644 on its own page on June 18 (last modified that day). On that same page the single number 86,644 is labeled three ways — **"compromised devices"** in the summary table, **"compromised access points"** in the body, and **"confirmed working login credentials"** in the FAQ — while the headline reads "80,000+." We report the dated values and the source's own labels, and let the reader weigh them.

For reference, the figure has not settled across the reporting. As of June 17, Hudson Rock cited more than 73,000 unique Fortinet URLs and SOCRadar more than 30,000 devices — already different units (URLs are not devices are not organizations). By June 19, SecurityWeek had noted SOCRadar's revision from 30,000 to 86,000, and CISA's June 18 advisory cited approximately 74,000 devices. Across these accounts the number shifts with each retelling rather than converging on a value — which is the measurement problem this analysis sets out to avoid.

METHOD

How we counted — stated before what we found

01

Everything below is reproducible by anyone holding comparable data. The method is fixed first; the results are reported second.

We measured a single quantity: the number of **distinct registrable organizations** whose remote-access portal credentials appear in our infostealer-log corpus. Not credentials, not URLs, not IP addresses, not devices — organizations, de-duplicated to the registrable domain. This is the only unit that supports a clean cross-vendor comparison, and it is deliberately the most conservative of the available counts.

- **Unit.** Distinct organizations (registrable domain). A single org with fifty captured credentials counts once.
- **Source.** Beacon Technology Group's CYFAX stealer-log corpus — credentials captured from compromised endpoints, aggregated from roughly 20,000 monitored sources across the criminal credential-trading ecosystem (channels, forums, and paste/dump venues). Underlying records are not published; the method below is fully reproducible against any comparable dataset.
- **Fingerprint.** Each vendor's remote-access portal has a near-unique login path. We matched on that path, anchored to the host to eliminate substring collisions,

and tallied per organization.

- **False positives.** The host-anchor (//[[^]/]+/...) removes coincidental substring matches (e.g. unrelated applications with a similar path). Records failing the anchor were excluded.
- **Reported exclusions.** SonicWall and Sophos were deliberately excluded: their portal paths lack a distinctive fingerprint and any count would be contaminated by false positives. We decline to publish a number we cannot defend — their absence here is a limit of the method, not evidence of safety.

CONTROL — A REPORTED NULL RESULT

We tested four of the named “victims” against the same corpus, including a known Fortinet customer. In that limited sample we did not observe internet-facing FortiGate SSL-VPN portals; the captured employee credentials resolved instead to federated identity (ADFS) and single-sign-on. Four organizations is a small sample and we draw no general conclusion from it — but it is a result that runs against the convenient story rather than toward it, and we report it for that reason: at least some of the marquee names do not appear to run the surface FortiBleed exploits.

4

named “victims” sampled, incl.
a known Fortinet customer

0

found running internet-
facing FortiGate SSL-VPN

100%

resolved to ADFS / SSO
portals — a different attack
surface

Reproducibility. The portal-path fingerprints used for each vendor are listed in the appendix. Anyone with a comparable credential-exposure dataset can apply the same match criteria and obtain a comparable count. That is the test of a measurement: someone else can run it.


The same exposure, across the category

One unit, one method, five vendors. The vendor in this week’s headlines lands toward the lower end of the range — fourth of five, and neither the most exposed nor the least.

EXHIBIT B

Distinct organizations with remote-access portal credentials in stealer-log data, by vendor

Each bar is distinct organizations (de-duplicated to registrable domain), same corpus, same date, same match method, across all five platforms.



Measured identically across five platforms, the counts run from 1,744 (Palo Alto) to 3,607 (Citrix). **Fortinet — the vendor in this week’s headlines — sits fourth of five at 2,799, with three platforms above it and one below.** The spread tracks how much remote-access surface each platform has put on the internet over the years, not how secure it is. The vendor that received a brand name, a victim list, and dedicated lookup portals this week is, by this measurement, toward the lower end of the set — neither the most exposed nor the least.

VENDOR & PLATFORM	PORTAL FINGERPRINT	DISTINCT ORGANIZATIONS
Citrix NetScaler / Gateway	/logon/LogonPoint	3,607
Ivanti Pulse Connect Secure	/dana-na/auth	2,996
Cisco ASA / AnyConnect	/+CSCOE+/ 	2,901
Fortinet FortiGate SSL-VPN	/remote/login	2,799
Palo Alto GlobalProtect	/global-protect/login	1,744

Same corpus, same date, same de-duplication to registrable domain. SonicWall and Sophos excluded for lack of a distinctive portal fingerprint.

READING THE RESULT

What this does — and does not — say

03

The discipline that makes a number trustworthy is stating its limits as plainly as its findings.

Exposure is not compromise. A captured portal credential means an endpoint somewhere leaked a login to that surface — not that the device is breached or the organization is owned. The honest claim is narrow: credentials for these portals circulate in commodity stealer-log data. It applies equally to every vendor in the table.

And the unit keeps changing. By June 19 the same campaign was being counted in at least five different ways: roughly 1.16 billion credential *attempts*; more than 320,000 *targeted devices*; 74,000 to 86,644 *compromised devices*, depending on the source; 86,644 working *credentials*; and, where one firm matched the leaked data against its own telemetry, 845 affected *organizations*. Each counts something

genuinely different, and a headline is free to reach for the largest. We count the most conservative of them — distinct organizations — for exactly that reason: it is the unit hardest to inflate.

The ordering is real, but it is not a security ranking. Citrix shows more exposure than Palo Alto in this corpus; that does not make Citrix “worse” or Palo Alto “safer.” A count like this rises with how many internet-facing portals a platform has in the field, and how long they have been there — deployment footprint and age, not the quality of the product behind them. Reading a league table of vendor security into these numbers would repeat the error this analysis exists to correct, simply pointed in a new direction.

The spread is shaped by deployment, not defect. Citrix NetScaler and the Cisco and Fortinet fleets are among the largest and oldest internet-facing remote-access footprints in the world; Palo Alto’s GlobalProtect is comparatively newer and more cloud-delivered. That the counts sort in roughly that order is the actual signal: **exposure tracks the category and its footprint, not the vendor.** The platform that drew the headline lands toward the lower end of the range — more exposed than one peer, less than three — which is unremarkable for any widely deployed platform, and not the profile of a singular outlier. A larger number for any vendor most plausibly means more surface deployed, and should never be read as a relative security judgment.

These are corpus-relative measurements. They describe representation within this credential-exposure corpus — not absolute estimates of global market exposure. We deliberately do not convert them into an “exposure rate” against deployed-device counts: reliable per-vendor deployment denominators do not exist, and dividing a firm number by a guessed one would manufacture exactly the kind of false precision this analysis is built to avoid. The raw, corpus-relative count is the honest figure, and we leave it as that.

Why the credential layer stays invisible

The reason this pattern persists — for every vendor — is that it lives in a layer the prevailing defensive model does not watch. Per-device, per-vulnerability scanning evaluates the appliance. It has nothing to say about a valid credential sitting in someone else’s database.

WHAT THE DEVICE SEES

A valid login. The appliance authenticates whoever presents correct credentials, exactly as designed. A vulnerability scanner reports it patched and green.

WHAT IT CANNOT SEE

That the same credential is already in an attacker's validated list — harvested off an endpoint, circulating in stealer logs, usable long after any underlying flaw was patched.

No CVE fires

No failed-login alert

Credential outlives the patch

The same blind spot explains the shape of the headline. The industry watches artifacts — individual devices, individual vulnerabilities — and rarely measures the credential layer between them, where exposure accumulates continuously and across every vendor at once. Point a spotlight at one slice of that layer and attach a name, and a standing condition reads as a sudden, five-figure event — its scale a function of how much was never being counted, and, as the figures earlier showed, of how loosely the count is defined. Branding one vendor changes none of that: not where the exposure lives, and not how long it has lived there.

ON MEASUREMENT

A note on standards, naming no one. Credible analysis states its unit, its method, and its limits, and reports the tests that failed to confirm its thesis. A figure that shifts severalfold within days — while blending devices, addresses, and credentials into a single number — is a headline, not a measurement. The difference is not tone; it is whether someone else can pick up the method and arrive at the same place.

What defenders should take from this is vendor-neutral, and the order matters — because against credentials harvested from infected endpoints, the controls are not equally effective. Every platform in this analysis already supports all of them; the exposure persists where they are not enforced, not where they are unavailable.

Multi-factor authentication, on every externally reachable gateway, is the control that defeats this dataset. A credential captured from an endpoint — however long, however recently rotated — does not authenticate without the second factor. It

should be the floor on every device, not an option. Next, keep management interfaces off the public internet, which removes the place a stolen credential would be replayed. Then monitor the credential layer — whether your credentials are already circulating — because once a credential reaches a stealer log, every preventive control upstream has already been passed. Re-authenticate administrators after firmware updates so stored credential hashes are re-derived under current protections. Strong, screened credentials remain sound hygiene, but they are a single line item, not the fix: an infostealer captures a password in plaintext as it is typed, so length and rotation do nothing once the endpoint is compromised.

This is, in substance, what the U.S. Cybersecurity and Infrastructure Security Agency recommended. CISA's June 18 advisory urges affected operators to reset credentials, store administrator hashes with PBKDF2, review authentication and domain-controller logs, enforce phishing-resistant MFA on every external gateway, and remove management access from the public internet — and it directs no one to patch a Fortinet product, because the activity it describes is credential exposure, not a software flaw. The official guidance and the measurement here point the same way: the fix lives in credential hygiene and configuration, which every platform in the table already supports.

The appliance being current is necessary. It was never sufficient — and that is true of every name in the table.

APPENDIX — FINGERPRINTS & REFERENCES

Portal-path fingerprints (host-anchored)

Fortinet FortiGate	<code>//[^\s]+/remote/login</code>
Cisco ASA / AnyConnect	<code>//[^\s]+\+CSCOE\+/</code>
Palo Alto GlobalProtect	<code>//[^\s]+/global-protect/login</code>
Ivanti Pulse Connect Secure	<code>//[^\s]+/dana-na/auth</code>
Citrix NetScaler / Gateway	<code>//[^\s]+/logon/LogonPoint</code>

Each pattern is host-anchored (`//[^\s]+/`) so the path must follow the host, not appear as a substring elsewhere in the URL. The Citrix pattern matches both the `/index` and `/tminindex` StoreFront variants; hosts on non-public namespaces (`.corp`, `.intra`, bare

IPs) do not resolve to a registrable organization and are excluded, making the Citrix count conservative. SonicWall and Sophos are excluded for lack of a single distinctive portal path.

References

- Arghire, I. (2026, June 19). *FortiBleed: 86,000 Fortinet device credentials compromised*. SecurityWeek. Retrieved June 19, 2026, from <https://www.securityweek.com/...>
- Beaumont, K. (2026, June 17). *FortiBleed — 75k Fortinet firewalls have admin passwords cracked*. DoublePulsar. Retrieved June 18, 2026, from <https://doublepulsar.com/>
- Cybersecurity and Infrastructure Security Agency. (2026, June 18). *CISA urges hardening Fortinet devices after reports of credential exposure* [Alert]. Retrieved June 19, 2026, from <https://www.cisa.gov/...>
- Franceschi-Bicchierai, L. (2026, June 17). *Cybercriminals allegedly hacked tens of thousands of Fortinet firewalls used by major companies all over the world*. TechCrunch. <https://techcrunch.com/2026/06/17/...>
- GreyNoise. (2025, December 17). *Coordinated credential-based campaign targets Cisco and Palo Alto Networks VPN gateways*. Retrieved June 18, 2026, from <https://www.greynoise.io/blog/...>
- Hudson Rock. (2026, June). *FortiBleed: 75,000 Fortinet firewalls compromised*. Infostealers. Retrieved June 18, 2026, from <https://www.infostealers.com/>
- Huntress. (2026, June). *2026 June — FortiBleed credential exposure*. Retrieved June 19, 2026, from <https://support.huntress.io/...>
- SOCRadar. (2026). *FortiBleed: The compromise of 80,000+ Fortinet firewalls* [Threat analysis]. (Original work published June 16, 2026; updated June 18, 2026.) Retrieved June 18, 2026, from <https://socradar.io/blog/fortibleed-fortinet-firewalls-compromised/>

Credential-exposure figures are derived from Beacon Technology Group’s CYFAX stealer-log corpus, aggregated from roughly 20,000 monitored sources across the criminal credential-trading ecosystem. Query method and match criteria are stated in §01 and above; underlying records are not published. Third-party figures were retrieved June 18–19, 2026, and reflect each cited source as published on those dates.

Measured, not branded.

This analysis was produced by the Beacon Technology Group HUMINT Team. Every figure traces to a stated source or a reproducible method. Where a number could not be measured cleanly, it was left out rather than estimated.

Beacon Technology Group · HUMINT Team · June 2026 · detect.solutions