EOTISEC
ANALYTICAL INTELLIGENCE REPORT

EOTISEC Analytical Division

Report Title	**FortiBleed: Mass Credential Compromise of Internet-Facing Fortinet FortiGate Firewalls and SSL VPN Gateways**
Report Number	**EOTISEC-2026-044**
Date of Report	**18 June 2026**
Coverage Period	**December 2025 — June 18, 2026**
Classification	**SENSITIVE BUSINESS DOCUMENT**
Originator	**EOTISEC Analytical Division**
Subject	**A threat actor has assembled a validated, searchable database of working administrator and SSL VPN credentials for internet-exposed Fortinet FortiGate devices across 194 countries. The campaign is active. The initial access vector is unconfirmed and no Fortinet zero-day has been confirmed.**
Customer Sector	**Information Technology and Communications (DHS Critical Infrastructure Sector). Secondary: Financial Services, Government Facilities, Energy, Defense Industrial Base**
Distribution	**Subscribers**

SECTION 1: SCOPE AND PURPOSE

This report assesses the credential compromise campaign tracked publicly as FortiBleed, disclosed in mid June 2026 and affecting internet-facing Fortinet FortiGate firewalls and SSL VPN gateways. The assessment is written for subscribers who operate, defend, or carry counterparty exposure to Fortinet perimeter devices. It draws on the original discovery by independent researcher Volodymyr Diachenko, corroborating analysis by Kevin Beaumont in collaboration with Hudson Rock, an independent infrastructure discovery by SOCRadar, vendor statements from Fortinet, and reporting by BleepingComputer, Help Net Security, CSO Online, Security Affairs, and Arctic Wolf. The report draws exclusively from open source reporting, government data releases, and market data.

Confidence in this product is uneven by design, and the reader should treat the three judgments below as carrying different evidentiary weight. The fact that a large, validated credential set exists and that the operation remains active rests on direct verification by multiple independent parties and is assessed with high confidence. The precise count of affected devices rests on figures that the disclosing parties themselves describe as not independently verifiable, and is assessed with moderate confidence. The method by which device configurations were first obtained, and the identity of the operators, rest on circumstantial indicators and competing expert readings, and are assessed with low confidence. No element of this report should be read as confirmation of a Fortinet zero-day vulnerability, because none has been confirmed as of the issue date.

SECTION 2: KEY JUDGMENTS

Judgment 1: Validated Credential Set and Active Campaign — CONFIRMED

We assess with **high confidence** that a threat actor has assembled a validated, searchable database of working administrator and SSL VPN credentials for internet-exposed Fortinet FortiGate devices, and that the operation was still running at the time of disclosure. Independent verification of a sample of live, working logins by a recognized researcher, combined with a separately discovered operational server holding the operators' tooling and victim database, supports this judgment.[1][2]

Judgment 2: Scale of the Compromise — QUANTIFIED WITH UNCERTAINTY

We assess with **moderate confidence** that the validated credential set covers on the order of 73,000 to 87,000 distinct devices, representing a substantial share of the internet-facing Fortinet population. The figure is uncertain because the disclosing parties report materially different counts and have stated the totals cannot be independently confirmed. The firmest floor is the roughly 30,791 credentials SOCRadar independently validated as working.[2][3][4]

Judgment 3: Initial Access Vector — UNRESOLVED

We assess with **low confidence** that the configuration files underpinning the credential set were obtained through a combination of known unpatched authentication-bypass vulnerabilities, recycled credentials from prior leaks and infostealer logs, and possibly an undisclosed access method. Expert readings diverge, and the access vector has not been established by any party.[3][5][6]

Judgment 4: Attribution — RUSSIAN-SPEAKING CRIMINAL GROUP, TENTATIVE

We assess with **moderate confidence** that the operators are a multi-operator, Russian-speaking criminal group, based on tooling, infrastructure choices, and a victim set weighted toward NATO member states. Attribution is among the weakest-evidenced elements and may shift as analysis continues.[1][7]

Judgment 5: Lateral Movement Risk — THE PRIMARY NEAR-TERM THREAT

We judge it **likely** that the most consequential near-term risk to affected subscribers is not the credential set in isolation but its use as a foothold for lateral movement into internal networks, including Active Directory environments, given the operators' documented pivot behavior and the privileged position of a compromised firewall.[5][8]

SECTION 3: SITUATION AND BACKGROUND

Fortinet FortiGate firewalls and SSL VPN gateways are among the most widely deployed network perimeter devices in the world, controlling access to internal networks across virtually every sector. A compromised FortiGate is therefore not a single-host problem. The device sits at the boundary of the protected network, holds credentials and routing configuration, and can be reconfigured by an administrator-level account to grant further access. This combination of ubiquity and privilege is the reason Fortinet perimeter devices have been a recurring target of both criminal and state-aligned operators over the past several years.

The FortiBleed dataset surfaced in mid June 2026. Independent researcher Volodymyr Diachenko reported finding an internet-exposed server holding what appeared to be valid Fortinet VPN credentials, including usernames, email addresses, and plaintext passwords, for organizations worldwide.[6] Kevin Beaumont independently reviewed portions of the data, obtained the wider set through Hudson Rock, and confirmed

that a sample of the administrator logins and passwords were authentic and that affected devices remained online and reachable.[3][5] Separately, SOCRadar reported discovering the operators' own infrastructure, including automation scripts, credential-testing tooling, logs, and a victim database organized by country, sector, and revenue.[1]

Why the Name Overstates the Mechanism

The label FortiBleed evokes Heartbleed and implies a single dramatic memory-disclosure flaw in the product. The evidence does not support that framing. SOCRadar stated plainly that it found no evidence Fortinet itself had been compromised or that a zero-day was involved, and assessed the credentials were most likely obtained through brute-force and credential-stuffing against internet-facing services.[1] Fortinet told reporters it was aware of a credential-theft campaign and that the data consisted of information from previous incidents together with credentials cracked through brute force.[4] The naming is doing rhetorical work that the confirmed facts do not. We retain the name FortiBleed throughout this report only because it is the common public identifier, not because it describes a discrete vulnerability.

LINCHPIN ASSUMPTION

This assessment assumes the credential set is what the disclosing parties represent it to be: a collection of working credentials harvested from device configurations, rather than a recycled aggregation of older leaks repackaged to appear new. ***If that assumption is wrong*** *and the bulk of the data is stale, as Fortinet's framing partly suggests and as occurred with the 2025 Belsen Group leak, the near-term risk to organizations that have rotated credentials since their last known exposure falls sharply, and the analytic line shifts from active compromise toward residual exposure. The independent verification of live, working logins by Beaumont is the single most important piece of evidence holding this assumption in place. If that verification were shown to rest on a small or unrepresentative sample, confidence in the scale judgment would degrade in step.[3][4][9]*

SECTION 4: ANALYSIS

The campaign combines three technical elements: extraction of device configuration files, offline cracking of credential hashes contained in those configurations, and a self-reinforcing harvesting loop that recycles recovered credentials back into the operation. Each is discussed below. The reader should note that the first element, how configurations were obtained, is the least understood and the point on which expert opinion diverges most sharply.

4.1 Evidence of Configuration-Level Exfiltration

The strongest technical signal in the dataset is that it contains information available only from inside a device configuration export, not from a login screen. Beaumont noted that the data includes internal email addresses and other elements visible only on the device itself, and stated that the data appears to have come from exports of configuration from the devices.[5][3] This distinguishes FortiBleed from a simple credential-scrape against a login interface and points to configuration-level access. One technical account describes the decryption path in concrete terms: FortiGate configuration files are encrypted, but the encryption key is derived from the device serial number, which is often visible on the management interface login page, so possession of the configuration file and the serial number makes decryption straightforward.[9]

4.2 The Hashing Weakness That Made Cracking Feasible

Older FortiOS versions stored administrator credentials using a legacy SHA-256 with salt scheme. SHA-256 is a fast hash, which is desirable for integrity checking and damaging for password storage, because speed is exactly what an offline cracker exploits.[3] Fortinet hardened credential storage in early 2025 by

migrating to PBKDF2 with a randomized salt, a deliberately slow scheme that resists offline cracking. Arctic Wolf documented the precise versions: PBKDF2-based hashing was introduced in FortiOS 7.2.11, 7.4.8, and 7.6.1, replacing the legacy SHA-256 storage.[10]

The protection had a critical gap. When a device was upgraded from an earlier version, existing administrator passwords remained stored as SHA-256 hashes until the corresponding administrator logged in again after the upgrade.[10] Many devices were upgraded but never had every administrator log back in, so they continued to store crackable SHA-256 hashes. Arctic Wolf further noted, citing Fortinet, that even after a password is updated to PBKDF2 the previous SHA-256 hash is retained in a hidden old-password setting for backward compatibility, which has its own residual-exposure implications.[10] Beaumont's analysis attributes the cracking to a distributed cluster, reported variously as 45 or 48 GPUs, managed through the open-source Hashtopolis framework, capable of pushing very large candidate volumes per second against fast hashes.[3][11]

Note on a Numeric Discrepancy

Public reporting cites the cracking cluster as both 45 GPUs and 48 GPUs across otherwise reliable outlets. The discrepancy is immaterial to the assessment, since either figure supports the same conclusion about offline cracking capability, but it is flagged here so the reader does not treat a single precise figure as settled.[3][11]

4.3 The Self-Reinforcing Harvesting Loop

SOCRadar's reconstruction from the operators' own infrastructure describes a fully automated cycle. The operators scan the internet for Fortinet devices, test each against a curated list of known passwords, and record every successful login. A compromised device is then used as a listening post to capture additional credentials passing through it, and those credentials are fed back into the scanner to compromise further devices, so the system feeds itself.[1] The password list is not random. It is assembled from credentials leaked in earlier Fortinet incidents, which works precisely because many organizations never rotated credentials after a prior breach.[1] Diachenko's reconstruction describes the operators intercepting SSL VPN authentication hashes, cracking them offline, and using the recovered passwords to pivot into internal Active Directory environments.[6][8]

4.4 Scale Figures and Why They Disagree

The reported scale varies by source, and the variation is itself analytically important. The table below records the principal figures as stated by each party, without reconciling them into a single number, because the disclosing parties have not reconciled them and have stated the totals cannot be independently verified.[2][3][4]

Source	Reported Figure	Basis as Stated
Hudson Rock / Beaumont	**73,932 URLs; ~75,000 devices**	Approx. 50% of internet-facing Fortinet firewalls indexed by Shodan at disclosure
SOCRadar	**86,644 devices; 80,000+ IPs**	Count of entries in the operators' own validated database
SOCRadar (verified)	**30,791 working credentials**	Subset SOCRadar independently validated as working
Arctic Wolf	**30,000 to 75,000 devices**	Range across the analyses it reviewed

Source	Reported Figure	Basis as Stated
Diachenko (attempts)	**~1.16 billion attempts**	Against 320,777 FortiGate targets, plus 2.1 billion against 163,650 MSSQL servers, per recovered logs

Two observations follow. First, the verified subset of roughly 30,791 working credentials is a firmer floor than the headline 73,000 to 87,000 device counts, and a cautious reader should anchor on the verified floor rather than the upper estimates.[1][2] Second, the campaign reached well beyond Fortinet: the same operators ran a parallel brute-force effort against more than 160,000 Microsoft SQL Server systems, which indicates a broad credential-access operation rather than a Fortinet-specific exploit.[6]

4.5 MITRE ATT&CK Mapping

The observed and reported behaviors map to the following ATT&CK techniques. The mapping reflects activity described by the disclosing parties. Techniques tied to the unconfirmed initial access vector are marked as assessed rather than confirmed.

Technique	Name	Tactic	Basis
T1595	Active Scanning	Recon	Confirmed
T1190	Exploit Public-Facing Application	Initial Access	Assessed
T1133	External Remote Services	Initial Access	Confirmed
T1110.002	Brute Force: Password Cracking	Credential Access	Confirmed
T1557	Adversary-in-the-Middle	Collection	Reported
T1003	OS Credential Dumping (config extraction)	Credential Access	Assessed
T1078	Valid Accounts	Initial Access	Confirmed
T1136	Create Account (backdoor admin)	Persistence	Reported
T1556	Modify Authentication Process	Defense Evasion	Assessed

On the defensive side, the relevant MITRE D3FEND countermeasures are credential rotation and revocation, multi-factor authentication enforcement, network traffic filtering to remove management-interface exposure, and inbound session analysis to detect anomalous administrative logins. These are reflected in the recommendations in Section 5.

4.6 The Initial Access Question

The unresolved question at the center of this campaign is how the operators first obtained device configurations at scale. The disclosing parties agree they do not know. BleepingComputer reported that none of Diachenko, Hudson Rock, or Beaumont had identified how the configuration data was originally obtained, and that it was unclear whether it came from previously disclosed vulnerabilities, a newly discovered flaw, or another method.[6] CSO Online recorded the same gap, quoting researchers that the initial access vector is presently unknown.[8] We therefore present the competing hypotheses rather than asserting a conclusion.

Hypothesis A: Known Unpatched Vulnerabilities

The leading candidate among disclosed flaws is CVE-2026-24858, a FortiCloud single sign-on authentication bypass. CISA, citing Fortinet, assigns it CWE-288 and a CVSS score of 9.4, and added it to the Known Exploited Vulnerabilities catalog on 27 January 2026 with a federal remediation deadline of 30 January 2026.[12][13] On devices with FortiCloud SSO enabled, an attacker with their own FortiCloud account and a registered device could authenticate to other customers' devices, create local administrator accounts for persistence, and make unauthorized configuration changes, behavior Arctic Wolf observed directly.[12][14] Critically, this flaw compromised devices that were fully patched against the December 2025 SSO bypasses CVE-2025-59718 and CVE-2025-59719, which means patch currency alone did not protect against it.[12][15] This hypothesis fits the observed configuration exfiltration, but it is constrained. FortiCloud SSO is not enabled by default, and the number of exposed SSO-enabled devices was in the low tens of thousands and falling after disclosure, which is smaller than the FortiBleed device count.[15]

Hypothesis B: An Undisclosed Vulnerability

Beaumont raised the possibility that the operators used a previously unknown flaw to obtain configurations, reasoning from the breadth of the data and the presence of devices on recent patches.[3][11] This hypothesis explains why the dataset includes devices that appear current, but it rests on inference from effect rather than on any identified flaw, and no such vulnerability had been confirmed by Fortinet as of the issue date.[3]

Hypothesis C: Recycled Credentials and Infostealer Logs

SOCRadar's position is that the credentials were most likely obtained through brute-force and credential-stuffing using credentials leaked in earlier incidents, with no Fortinet zero-day involved.[1] Fortinet's own statement matches this in part, describing the data as a mix of information from previous incidents and brute-forced credentials.[4] This hypothesis is the most parsimonious and best explains the parallel MSSQL campaign, but it sits in tension with the configuration-level artefacts Beaumont identified, which credential-stuffing alone would not produce.[5]

The most defensible reading, and the one we adopt at low confidence, is convergence rather than a single cause: known unpatched flaws on some devices, an undisclosed method on others, and recycled credentials across the rest, with offline cracking of weak hashes tying the set together. We flag explicitly that this convergence reading is partly a function of incomplete evidence, and that it should not harden into a settled conclusion as further analysis emerges.

4.7 Historical Pattern of Comparable Activity

FortiBleed is not an isolated event but the latest in a multi-year pattern of credential and configuration compromise targeting Fortinet perimeter devices. Placing it against prior incidents both calibrates the scale claims and clarifies what is and is not novel about it.

Belsen Group, January 2025

In January 2025 a threat actor calling itself Belsen Group released full configuration files and VPN credentials for more than 15,000 FortiGate devices, organized by country and IP address, free of charge on a criminal forum.[16][17] Beaumont and CloudSEK assessed the data had been collected in October 2022 through exploitation of the authentication-bypass flaw CVE-2022-40684 while it was a zero-day, with SSL VPN credentials sourced via the older path-traversal flaw CVE-2018-13379.[17][18] Fortinet confirmed the configurations were genuine but characterized the release as dated 2022 data aggregated to appear new, noting most affected devices had long since been upgraded.[19] Censys found that as of mid-January 2025 over half the 15,469 compromised hosts were still online and roughly a third still exposed their web login interfaces, which is why aged configuration data remained dangerous.[20] Beaumont states the FortiBleed IP addresses are largely different from the Belsen set, indicating a more recent and larger collection rather than a rerelease.[3][5]

The 2021 SSL VPN Credential Dump

In September 2021 a threat actor published SSL VPN credentials for nearly half a million Fortinet accounts, harvested by exploiting the same CVE-2018-13379 path-traversal flaw later implicated in the Belsen leak.[17] That incident established the template that recurs through FortiBleed: a known, patchable flaw, a long tail of unpatched or unrotated devices, and credentials that retain value for years because organizations fail to rotate them after exposure.

The December 2024 Management-Interface Campaign

Arctic Wolf documented a campaign beginning in early December 2024 against FortiGate devices with management interfaces exposed on the public internet. At the time the point of intrusion was not tied to a specific CVE.[19] This episode prefigures the central FortiBleed risk factor, namely management-interface exposure, which Beaumont reports characterizes a majority of the FortiBleed-affected devices.[5]

The FortiCloud SSO Bypass Chain, December 2025 to January 2026

The immediately preceding episode is the FortiCloud SSO bypass chain. CVE-2025-59718 and CVE-2025-59719, disclosed in December 2025, allowed unauthenticated attackers to bypass SSO via crafted SAML messages, and observed attacks involved authenticating as admin and immediately downloading the system configuration file, which contains hashed credentials.[15][21] When patched devices continued to be breached in January 2026, Fortinet identified the net-new CVE-2026-24858, disabled FortiCloud SSO service-wide on 26 January 2026, and issued its advisory on 27 January 2026.[12][13][22] This chain is the strongest historical link to FortiBleed because its documented outcome, attacker-driven download of configuration files containing hashed credentials, is exactly the artefact FortiBleed monetized at scale.[21]

4.8 Impact Assessment

The impact of FortiBleed is best understood not as the value of the credentials themselves but as the access those credentials unlock. Beaumont states the practical effect directly: with a working credential an attacker can log in remotely and gain access to the firewall and therefore the network behind it, change settings including security controls, and create backdoor administrator accounts.[8] A compromised perimeter device is a position from which to move laterally, not an endpoint of the attack.

Sectoral and Geographic Exposure

SOCRadar's analysis of the victim database found exposure across every major sector, with telecommunications the most heavily represented by volume at 5,616 entries, a concern that compounds because telecom infrastructure underpins communications for other sectors.[1] Government exposure totalled 591 entries across 111 domains, with India accounting for over 60% of government entries, and the dataset also reaches banks, hospitals, universities, energy companies, and large multinationals.[1] Geographically, India and the United States together account for nearly a third of all entries, with the remainder spread across Asia, Europe, the Americas, the Middle East, and Africa.[1] Enterprises above one billion dollars in revenue account for over 20% of entries.[1]

The Lateral Movement Risk

The documented pivot from FortiGate credentials into internal Active Directory environments is the mechanism by which a perimeter compromise becomes a domain compromise.[6][8] Where a firewall account is reused for, or trusted by, internal directory services, recovery of the firewall credential can seed privilege escalation inside the network. The risk is highest for organizations that operate flat networks, reuse administrator credentials across perimeter and internal systems, or have not segmented management access.

The Reported Defense-Sector Exposure

Diachenko reported that at least four organizations across Japan, Taiwan or Vietnam, Iraq, and Turkey were fully compromised, including a Turkish NATO defense contractor from which classified defense documents were allegedly exfiltrated, and SOCRadar reported recovering credentials for what appeared to be a defense-industry VPN endpoint.[1][24] These specific claims are single-sourced to the disclosing researchers and are reported here as claims rather than established fact. If substantiated, they would raise the campaign's significance from large-scale criminal credential harvesting toward potential espionage value, which is part of why the attribution question matters. We assess these claims at low confidence pending corroboration.

The Initial-Access-Broker Indicator

A structural feature of the dataset points to its intended use. Beaumont and others note that each entry carries the organization's industry, revenue, employee count, and country, formatted in the manner common to initial access listings sold in criminal markets.[5][6] This suggests the set was assembled for sale or for coordinated deployment across a team rather than for a single operator's use. At the time of SOCRadar's publication the specific dataset had not yet appeared for sale on criminal forums, which gives affected organizations a closing window to act before wider distribution.[1]

4.9 The Linchpin Assumption and Alternative Readings

The load-bearing assumption is identified in the callout in Section 3 and is not repeated here. The principal alternative reading worth weighing against the active-compromise framing is Fortinet's: that the data is substantially dated and aggregated to appear new, mirroring the vendor's 2025 Belsen characterization. The evidence that holds the active-compromise reading in place is Beaumont's verification of live working logins on currently reachable devices and the largely novel IP set relative to Belsen.[3][5][19] We weigh both, retain the active-compromise reading at the confidence stated in the Key Judgments, and treat the dated-data alternative as the most credible competing hypothesis rather than as refuted.

SECTION 5: DECISION SUPPORT AND RECOMMENDATIONS

The following actions reflect converging guidance from CISA, Arctic Wolf, SOCRadar, Fortinet, and independent researchers. They apply regardless of whether an undisclosed vulnerability is later confirmed, because the known attack surface alone justifies them. Any subscriber operating an internet-facing FortiGate firewall or SSL VPN gateway should treat the device as potentially affected until proven otherwise.[1][8][10]

Immediate Actions

- Rotate all administrator and SSL VPN credentials now, prioritizing internet-exposed devices and any device involved in a previous credential exposure. Password complexity offers no protection when a credential has already been recovered in plaintext, so rotation, not strengthening, is the operative control.[1][10]
- Remove the FortiGate management interface from direct public internet exposure and restrict it to trusted internal networks. Management-interface exposure characterizes a majority of affected devices and is the single most effective reduction in attack surface.[5][10]
- Enforce multi-factor authentication on every administrative and remote-access account, so that a recovered password alone does not grant access.[8][10]

Configuration and Patch Hygiene

- Upgrade to a supported FortiOS version, then require every administrator to log in at least once after the upgrade to force re-hashing to PBKDF2. The migration does not take effect for an administrator until that administrator logs in.[10][8]

- Where a re-login is not feasible, manually reset remaining administrator passwords using a super_admin account to trigger PBKDF2 hashing.[8][10]
- Disable FortiCloud SSO administrative login unless it is operationally required, since CVE-2026-24858 compromised even fully patched devices where SSO was enabled. Verify the Allow administrative login using FortiCloud SSO setting.[12][15]

Detection and Response

- Review device login history for unfamiliar access by time, source location, or account, and audit the configuration for unauthorized local administrator accounts or VPN changes consistent with the observed CVE-2026-24858 behavior.[1][12]
- Where indicators of compromise are present, treat the device as breached: restore configuration from a known-clean version, rotate all connected credentials including any LDAP or Active Directory accounts the device touches, and engage incident response to assess lateral movement into internal networks.[8][14]
- Check the organization's domains against the free lookup tools published by Hudson Rock and SOCRadar to determine whether the organization appears in the dataset, treating any appearance as a strong indicator of compromise rather than a benign listing.[1][6]

Decision Support

The most time-sensitive decision for any subscriber operating internet-facing Fortinet devices is credential rotation paired with management-interface withdrawal, executed in the next several days rather than queued behind a patch cycle, because the credential set is already validated and the cost of inaction is unauthorized access that may already be in progress. This decision belongs with the Chief Information Security Officer. For financial-services and government subscribers carrying counterparty exposure to affected vendors, the time-sensitive decision is third-party assurance: confirming that material Fortinet-dependent vendors have rotated credentials and withdrawn management exposure. The window for first-mover containment closes as the dataset moves toward wider criminal distribution.[1][5]

SECTION 6: INFORMATION GAPS AND COLLECTION REQUIREMENTS

Four intelligence gaps would change this assessment if closed. We will revise this product as they close.

- The initial access vector. Confirmation or exclusion of an undisclosed vulnerability would reshape the patching guidance. Collection priority: high.
- Firmer attribution. Clarifying whether espionage value sits alongside the criminal motive bears on the significance of the defense-sector claims. Collection priority: moderate.
- Corroboration of the single-sourced defense-sector compromise claims, which currently rest on the disclosing researchers alone. Collection priority: high.
- A reconciled device count, which would resolve the scale uncertainty between the verified floor near 30,791 and the upper estimates near 87,000. Collection priority: moderate.

SECTION 7: SOURCE SUMMARY STATEMENT

This assessment rests on a mix of primary disclosures, a vendor statement, and reputable secondary security reporting, listed in the endnotes. Source reliability is uneven and is treated as such throughout. The discovery and verification claims by Beaumont and the infrastructure discovery by SOCRadar are first-order evidence from parties with direct access to the data and a track record in this domain, and carry the most weight. The CISA alert and Fortinet advisory on CVE-2026-24858 are primary sources at the highest

credibility level for the vulnerability detail. Vendor statements from Fortinet on the FortiBleed dataset itself are weighed as those of an interested party with strong incentives to characterize the data as dated. Scale figures originating from the disclosing researchers are reported as stated and not independently confirmed by EOTISEC. Single-sourced claims, including the defense-sector compromise, are flagged in text as claims.

Source limitations: the initial access vector is unestablished and the device counts are not independently verifiable. The CVSS score for CVE-2026-24858 was reported inconsistently in secondary coverage. A claim of up to 9.8 appears in at least one outlet and is inconsistent with the 9.4 score published by CISA and Fortinet. This report uses the authoritative 9.4 throughout. One numeric discrepancy in the underlying reporting, the 45-versus-48 GPU cracking cluster, is carried as flagged rather than resolved.

SECTION 8: ANALYTIC TRADECRAFT SELF-CERTIFICATION

Requirement	Status
Objectivity: free of advocacy, personal preference, or policy bias	CONFIRMED
Independence of political consideration: no judgment shaped to support a particular outcome	CONFIRMED
Timeliness: source material is current and the report is actionable by the customer	CONFIRMED. Source material current to June 18, 2026
Based on all available sources: gaps documented in Section 6	CONFIRMED
Source credibility described: source quality addressed in Section 7	CONFIRMED
Uncertainty expressed: probability and confidence language conforms to ICD 203 standards	CONFIRMED. Confidence levels stated and differentiated across judgments
Assumptions distinguished from facts: linchpin assumption identified	CONFIRMED. One LINCHPIN ASSUMPTION identified in Section 3
Alternatives incorporated: competing hypotheses addressed	CONFIRMED. Three hypotheses for the access vector in Section 4.6, plus the dated-data alternative in Section 4.9
Customer relevance addressed: business impact stated by category	CONFIRMED. Impact in Section 4.8, with recommendations in Section 5
Decision support included: time-sensitive actions, ownership, and cost of inaction	CONFIRMED. Section 5 decision support paragraph
Clear and logical argumentation: main message stated up front in Key Judgments	CONFIRMED
Source verification: primary-source or multi-link corroboration	CONFIRMED. One secondary CVSS error corrected to the CISA value

SECTION 9: ENDNOTES — VERIFIED SOURCE CHAIN

1. SOCRadar. *FortiBleed: The Compromise of 80,000+ Fortinet Firewalls*. Published 16 June 2026, updated 18 June 2026. High credibility. Primary discovery of the operators' infrastructure. https://socradar.io/blog/fortibleed-fortinet-firewalls-compromised/

2. Arctic Wolf. *Active FortiBleed Campaign Impacting Fortinet Devices Across 194 Countries*. High credibility. https://arcticwolf.com/resources/blog/active-fortibleed-campaign-impacting-fortinet-devices-across-194-countries/
3. BleepingComputer. *FortiBleed leak exposes Fortinet VPN credentials for 73,000 devices*. 18 June 2026. High credibility. Beaumont verification quotes and the unknown-vector statement. https://www.bleepingcomputer.com/news/security/fortibleed-leak-exposes-fortinet-vpn-credentials-for-73-000-devices/
4. heise online. *Massive attack on Fortinet firewalls? 74,000 devices affected by FortiBleed*. 18 June 2026. High credibility. Carries the Fortinet statement to TechCrunch. https://www.heise.de/en/news/Massive-attack-on-Fortinet-firewalls-74-000-devices-affected-by-FortiBleed-11336418.html
5. Help Net Security. *74,000 Fortinet firewall credentials exposed in FortiBleed data leak*. 18 June 2026. High credibility. Direct Beaumont quotes on config-export origin. https://www.helpnetsecurity.com/2026/06/18/fortinet-fortibleed-data-leak/
6. BleepingComputer, as in endnote 3, for Diachenko's LinkedIn disclosure, the 1.16 billion FortiGate attempts, and the parallel 2.1 billion MSSQL attempts. https://www.bleepingcomputer.com/news/security/fortibleed-leak-exposes-fortinet-vpn-credentials-for-73-000-devices/
7. CSO Online. *FortiBleed campaign exposes 75,000 Fortinet firewalls worldwide*. High credibility. SOCRadar attribution language and watchTowr commentary. https://www.csoonline.com/article/4186790/fortibleed-campaign-exposes-75000-fortinet-firewalls-worldwide.html
8. CSO Online, as in endnote 7, for Beaumont on impact and the unknown initial access vector. https://www.csoonline.com/article/4186790/fortibleed-campaign-exposes-75000-fortinet-firewalls-worldwide.html
9. Security Affairs. *FortiBleed Exposes Admin Passwords for 75,000 Fortinet Firewalls*. 18 June 2026. High credibility. Config decryption via serial number and the Belsen IP comparison. https://securityaffairs.com/193817/hacking/fortibleed-exposes-admin-passwords-for-75000-fortinet-firewalls.html
10. Arctic Wolf, as in endnote 2, for FortiOS PBKDF2 versions 7.2.11, 7.4.8, 7.6.1, the post-upgrade re-login requirement, and the retained old-password hash. https://arcticwolf.com/resources/blog/active-fortibleed-campaign-impacting-fortinet-devices-across-194-countries/
11. Kevin Beaumont, DoublePulsar. *FortiBleed: 75k Fortinet firewalls have admin passwords cracked*. 17 June 2026. Primary analysis, cited via corroborating reporting where the original was not directly retrievable. https://doublepulsar.com/fortibleed-75k-fortinet-firewalls-have-admin-passwords-cracked-60299faa65f8
12. CISA. *Fortinet Releases Guidance to Address Ongoing Exploitation of Authentication Bypass Vulnerability CVE-2026-24858*. 28 January 2026. Highest credibility. CWE-288, CVSS 9.4, KEV listing. https://www.cisa.gov/news-events/alerts/2026/01/28/fortinet-releases-guidance-address-ongoing-exploitation-authentication-bypass-vulnerability-cve-2026
13. The Hacker News. *Fortinet Patches CVE-2026-24858 After Active FortiOS SSO Exploitation Detected*. Confirms the 27 January 2026 advisory and 30 January federal deadline. https://thehackernews.com/2026/01/fortinet-patches-cve-2026-24858-after.html
14. Triskele Labs. *Critical Fortinet FortiCloud SSO Authentication Bypass Under Active Exploitation CVE-2026-24858*. On observed configuration theft and persistence behavior. https://www.triskelelabs.com/blog/critical-fortinet-forticloud-sso-authentication-bypass-under-active-exploitation-cve-2026-24858
15. CyberDesserts. *CVE-2026-24858: The Fortinet Patch That Wasn't*. On patched devices compromised and the Shadowserver SSO-enabled device count falling below 10,000. https://blog.cyberdesserts.com/cve-2026-24858-fortinet-sso-bypass/

16. Help Net Security. *Configuration files for 15,000 Fortinet firewalls leaked*. 16 January 2025 (Belsen Group). https://www.helpnetsecurity.com/2025/01/16/leaked-fortinet-fortigate-configs-vpn-credentials-ip-list/
17. Dark Reading. *15K Fortinet Device Configs Leaked to the Dark Web*. 17 January 2025. On CVE-2022-40684 and CVE-2018-13379 as the Belsen sources. https://www.darkreading.com/endpoint-security/15k-fortinet-device-configs-leaked-dark-web
18. SecurityWeek. *Data From 15,000 Fortinet Firewalls Leaked by Hackers*. 16 January 2025. On the October 2022 collection date and the CVE-2022-40684 zero-day timeline. https://www.securityweek.com/data-from-15000-fortinet-firewalls-leaked-by-hackers/
19. The Register. *Fortinet: FortiGate config leaks are genuine but misleading*. 17 January 2025. On Fortinet's dated-data framing and the December 2024 Arctic Wolf campaign. https://www.theregister.com/2025/01/17/fortinet_fortigate_config_leaks/
20. Censys. *Massive FortiGate Config Leak: Assessing the Impact*. On 15,469 hosts with over half still online and the 2021 leak of nearly 500,000 VPN credentials. https://censys.com/blog/fortigate-config-leak-impact/
21. Rapid7. *Critical Vulnerabilities in Fortinet CVE-2025-59718, CVE-2025-59719 Exploited in the Wild*. On admin authentication followed by immediate configuration download. https://www.rapid7.com/blog/post/etr-critical-vulnerabilities-in-fortinet-cve-2025-59718-cve-2025-59719-exploited-in-the-wild/
22. SOC Prime. *CVE-2026-24858: FortiOS SSO Zero-Day Exploited in the Wild*. On the two malicious FortiCloud accounts blocked 22 January and SSO suspension 26 January 2026. https://socprime.com/blog/cve-2026-24858-vulnerability/
23. MITRE ATT&CK, External Remote Services (T1133), citing NSA, CISA, FBI, NCSC. *Russian GRU Conducting Global Brute Force Campaign to Compromise Enterprise and Cloud Environments*. July 2021. https://attack.mitre.org/techniques/T1133/
24. The CyberSec Guru. *FortiBleed: How 75,000 Fortinet Firewalls Were Silently Compromised in 2026*. On the reported four fully compromised organizations and the Turkish NATO contractor claim. https://thecybersecguru.com/news/fortibleed-fortinet-firewall-credential-leak/