CSAI Foundation | Cloud Security Alliance

FortiBleed: Default Credential Exploitation and Mass Fortinet Compromise

How 86,000 Firewalls Were Cracked Through Unrotated Factory Credentials and Legacy Password Hashing

2026-06-20

AI-assisted Rapid Research

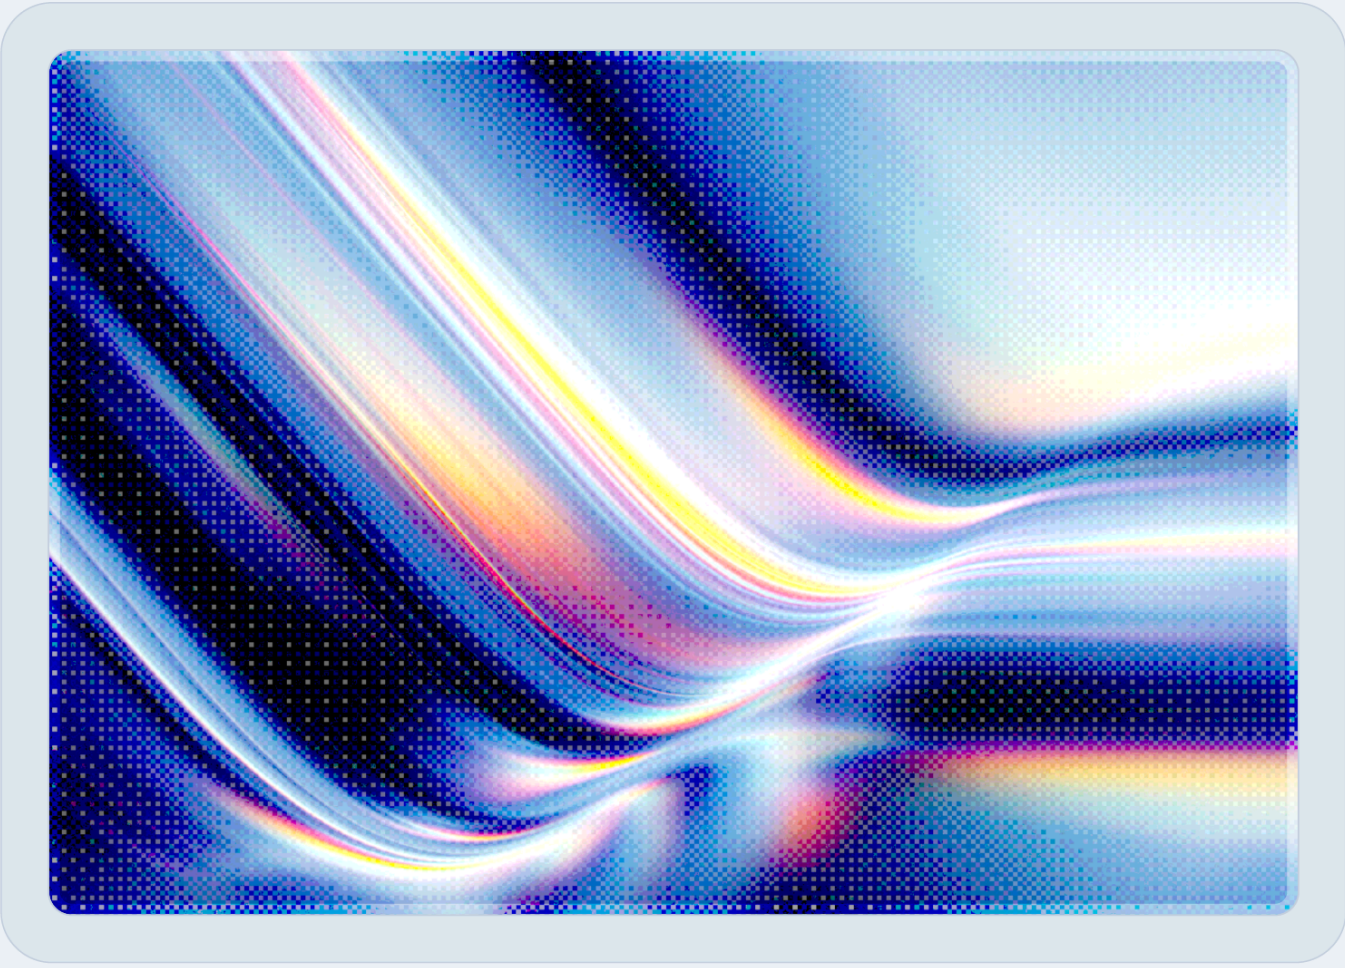

CSAI

CSA cloud security alliance®

Key Takeaways

- A large-scale credential compromise campaign, dubbed FortiBleed, has exposed verified administrator credentials for internet-facing Fortinet FortiGate firewalls across 194 countries — with reported device counts ranging from approximately 74,000 [3] to 86,644 [1] depending on publication date and verification stage — representing roughly half of all internet-reachable FortiGate devices worldwide based on Shodan data [2].
- Threat actors systematically extracted configuration backup files from vulnerable devices, then cracked the stored password hashes offline — a process made technically feasible at scale because many devices still stored credentials as salted SHA-256 hashes rather than the stronger PBKDF2 algorithm that Fortinet introduced in late 2025 [3].
- The majority of compromised accounts — approximately 63% — were either default Fortinet system accounts or generic administrator accounts that had never been renamed, enabling attackers to target the most predictable credentials before any brute-force effort was required [1].
- Separately, CVE-2026-24858, a critical FortiCloud SSO authentication bypass (CVSS 9.4), was exploited in the wild beginning in January 2026 and allowed attackers to create rogue local administrator accounts on fully patched devices [4][5].
- CISA issued an emergency advisory on June 18, 2026, ordering immediate session termination, credential rotation, MFA enforcement, and removal of FortiGate management interfaces from the public internet [6].
- Post-authentication activity by threat actors has escalated beyond firewall access into Active Directory compromise, lateral movement across internal segments, and persistent access to critical infrastructure in telecommunications, government, healthcare, finance, and energy sectors [1][7].

Background

Fortinet's FortiGate platform is widely deployed across enterprise data centers, government networks, healthcare systems, operational technology environments, and managed service provider (MSP) infrastructure worldwide, with internet-exposed instances numbering in the hundreds of thousands

according to Shodan scan data [2]. Its use spans a diverse range of critical sectors and geographies, making FortiGate a consistently high-priority target: an attacker who gains administrative credentials to a perimeter firewall inherits visibility into the internal network topology, the ability to modify routing and filtering rules, access to VPN tunnel endpoints, and — in many configurations — an authenticated foothold from which to pivot into Active Directory and other internal systems.

The FortiBleed campaign was formally disclosed on June 16, 2026, following independent analysis by SOCRadar and Arctic Wolf, who identified exposed operational infrastructure belonging to the threat actors [1][2]. SOCRadar's researchers discovered databases of validated credentials organized by country, sector, and target organization revenue, indicating a systematic and commercially oriented operation rather than opportunistic scanning. Reported device counts range from approximately 74,000 [3] to 86,644 [1] reflecting different publication dates and verification stages of the dataset. CISA issued its own advisory two days later on June 18, 2026, confirming active exploitation and issuing hardening guidance to affected organizations [6].

The campaign did not emerge in isolation. Beginning in late 2025 and accelerating through the first half of 2026, Fortinet devices were subjected to several overlapping exploitation efforts. In January 2026, a critical authentication bypass vulnerability designated CVE-2026-24858 was exploited against fully patched FortiOS devices, enabling unauthorized admin account creation through a flaw in FortiCloud's SSO implementation [4][5]. In parallel, a separate campaign observed by Amazon Threat Intelligence between January 11 and February 18, 2026, used AI-assisted tooling to enumerate exposed management ports and attempt credential-based access without exploiting any known CVE — relying entirely on weak passwords and single-factor authentication [8]. FortiBleed represents the convergence of these threads: large-scale automated exploitation, legacy credential storage weaknesses, and the enduring failure of many organizations to rotate default accounts.

Security Analysis

The Credential Harvesting Mechanism

The technical mechanics of FortiBleed center on a chain that connects configuration backup file access to offline credential cracking. FortiGate devices store administrator credentials — including their password hashes — in the device configuration. When an attacker can retrieve this configuration file, either through an unauthenticated vulnerability, an exposed management interface, or a credential obtained through an earlier stage of access, they obtain not just network topology information but the raw material to derive working passwords [3].

Fortinet addressed its password storage weakness by migrating from salted SHA-256 to the more computationally expensive PBKDF2 algorithm in FortiOS versions 7.2.11, 7.4.8, and 7.6.1, released in late 2025 [9]. However, this migration carried a critical conditional: the hash upgrade only applied to accounts whose owners actively logged in after upgrading firmware. Accounts that existed prior to the update and whose passwords were never subsequently changed remained stored as salted SHA-256 hashes, which are readily crackable offline using GPU-accelerated tooling. Compounding this, a backward-compatibility mechanism preserved previous SHA-256 hashes in a hidden `old-password` configuration field that, while not visible through the management interface, is fully accessible in any configuration backup taken by a super_admin account [3]. Devices believed to be protected by PBKDF2 could therefore still yield crackable hashes to an attacker with access to a configuration backup.

Once configuration files were in hand, the threat actors ran extracted hashes through GPU clusters. Published benchmark data indicates that a single modern GPU can test hundreds of millions of SHA-256 iterations per second, making passwords below adequate length or character-set diversity recoverable within hours to days. The verified, working credentials that resulted were then organized into structured databases segmented by target geography, sector, and estimated organizational value [1].

The Role of Default and Generic Accounts

A defining characteristic of the FortiBleed dataset is its composition by account type. SOCRadar's analysis found that generic administrator accounts accounted for approximately 35% of compromised credentials, while built-in Fortinet system accounts accounted for 28.3%, and organization-specific accounts made up the remaining 36.7% [1]. The implication is significant: more than six in ten compromised accounts had either never been renamed from factory defaults or used predictable generic names that threat actors could target with high confidence before investing any brute-force capacity. This is not a vulnerability in the conventional sense — it is a configuration hygiene failure that has persisted across tens of thousands of production deployments.

Default credentials represent a known and documented risk. Fortinet's own hardening guidelines recommend renaming or disabling the default `admin` account and removing built-in system accounts that are not operationally necessary. The prevalence of these accounts in the FortiBleed dataset suggests that configuration hardening checklists, even where formally required by policy, are inconsistently applied in practice — particularly in large or distributed environments where firewalls are deployed by different teams, maintained through managed service contracts, or inherited through acquisitions.

CVE-2026-24858: Authentication Bypass Through Trusted Federation

Separate from the credential cracking campaign, CVE-2026-24858 introduced a distinct class of risk affecting organizations that had enabled FortiCloud SSO for administrative access. The vulnerability carries a CVSS base score of 9.4 and is classified under CWE-288 (Authentication Bypass Using an Alternate Path or Channel) [4][5]. An attacker who controlled any FortiCloud account and a registered FortiGate device could exploit a cross-tenant authentication flaw to authenticate to other customers' devices as long as FortiCloud SSO was enabled on those devices — effectively using their own valid identity to assume administrative authority over unrelated organizations' infrastructure.

Active exploitation was confirmed in January 2026, when multiple Fortinet customers reported rogue local administrator accounts created on their firewalls despite running fully patched FortiOS versions [5]. CISA added CVE-2026-24858 to the Known Exploited Vulnerabilities catalog on January 27, 2026, with a remediation deadline of January 30, 2026 [6]. The vulnerability illustrates a pattern that appears repeatedly across enterprise security infrastructure: federated authentication features introduced for operational convenience can create trust relationships that extend an organization's attack surface in ways that are difficult to fully characterize and monitor.

AI-Assisted Attack Automation

A February 2026 investigation by Amazon Threat Intelligence documented a concurrent campaign against FortiGate devices that, while distinct from FortiBleed, reflects the broader industrialization of perimeter exploitation [8]. Between January 11 and February 18, 2026, a threat actor targeting FortiGate management ports used multiple commercial generative AI tools to accelerate attack development: generating targeted credential wordlists, automating reconnaissance scripts, synthesizing exploitation chains, and producing operator documentation. The campaign succeeded purely through credential abuse against exposed management interfaces — no CVE exploitation occurred. The integration of AI tooling appears to have reduced the expertise barrier for this kind of sustained, automated exploitation and likely increased the speed at which the actor could iterate on attack variations.

This campaign, combined with FortiBleed's GPU-accelerated hash cracking and the billions of authentication attempts against FortiGate targets during the same period [7], indicates that perimeter device exploitation has reached what Fortinet's own 2026 Global Threat Landscape Report terms "industrial scale" — machine-speed attacks operating continuously against a static population of internet-exposed management interfaces [10].

Post-Access Exploitation Chain

Administrator access to a FortiGate firewall is rarely the final objective. The credential dataset assembled in FortiBleed has been used to establish beachheads for follow-on operations that extend deep into the internal environments behind those firewalls. Confirmed post-exploitation activity has included creation of persistent VPN tunnels under legitimate-looking names, extraction of internal routing tables and network diagrams, and direct authentication against downstream infrastructure including Active Directory domain controllers, RADIUS servers, and internal management systems [7]. Full network compromises — where attackers established persistence and conducted lateral movement beyond the initial Fortinet device — have been confirmed at organizations across multiple countries [7].

The threat actor behind FortiBleed appears to be Russian-speaking based on operational infrastructure analysis, though this assessment is based on indirect indicators and formal attribution remains unconfirmed [1]. Regardless of origin, the operational capability demonstrated — GPU-scale offline cracking, billions of authentication attempts, and structured post-access data collection — is consistent with a well-resourced and professionally organized adversary rather than an opportunistic actor.

Sector exposure in the confirmed compromise data spans telecommunications (the most heavily represented sector), government entities across 111 domains, healthcare, financial services, education, and energy [1][7]. India and the United States together account for nearly one-third of identified compromises by device count [1], suggesting this may reflect the concentration of internet-exposed FortiGate deployments in those markets rather than any specific targeting priority.

Recommendations

Immediate Actions

Organizations running internet-facing Fortinet devices should treat the following as emergency response steps rather than scheduled maintenance, given that credentials from affected organizations may already be in active use by threat actors.

All active SSL VPN and administrative sessions should be terminated immediately to invalidate any sessions established with compromised credentials. Following session termination, all FortiGate administrator passwords must be rotated — particularly for any account using a default name such as `admin`, any built-in system account, or any generic account name that follows predictable patterns.

Password rotation is ineffective if the new credentials are reused across devices or set to values that can be recovered from adjacent systems, so rotation must be accompanied by enforcing unique, high-entropy passwords per device.

Organizations should audit their FortiOS version against Fortinet's PBKDF2 migration threshold (FortiOS 7.2.11, 7.4.8, and 7.6.1) and verify, after logging in to each device post-update, that administrator credentials are being stored using PBKDF2 rather than SHA-256. Consulting the Fortinet community guidance on the `old-password` hidden field is recommended to confirm that legacy hashes have been cleared from configuration backups [9]. Any organization that takes configuration backups should treat existing backup files as potentially sensitive and restrict their storage and access accordingly.

FortiCloud SSO should be disabled on all devices that have not yet patched to a version not affected by CVE-2026-24858. Even on patched devices, SSO should only be re-enabled after a deliberate security review that accounts for the cross-tenant trust implications of federated administrative access [4].

Short-Term Mitigations

Management interfaces for FortiGate devices — including the web management console, SSH, and FortiCloud access — should be removed from public internet exposure and restricted to trusted administrative networks, bastion hosts, or private-channel access paths such as out-of-band management networks [6]. CISA's advisory is explicit on this point: internet-exposed management interfaces are a precondition for the initial access phase of campaigns like FortiBleed, and removing that exposure eliminates an entire class of attack surface regardless of what vulnerabilities are later disclosed.

Phishing-resistant multi-factor authentication (MFA), using FIDO2 security keys or certificate-based authentication, should be enforced across all remote access paths and administrative interfaces [6]. Phishing-resistant MFA significantly raises the cost of credential-based attacks, including hash cracking, credential stuffing, and brute force. Weaker forms of MFA — SMS OTP and push notifications — remain vulnerable to adversary-in-the-middle proxies and push-fatigue attacks; FIDO2 or certificate-based implementations are required to fully eliminate this attack vector. For SSL VPN deployments specifically, CISA has mandated that all FortiGate SSL VPN deployments migrate to certificate-based or FIDO2 authentication.

Logs from FortiGate devices, SSL VPN concentrators, authentication systems, and domain controllers should be collected and reviewed for signs of unauthorized access — specifically looking for new account creation, unexpected administrative configuration changes, VPN sessions from unusual geographic origins, and authentication events that do not correspond to known administrator activity. Where a security information and event management (SIEM) system is available, detection rules aligned with FortiBleed post-exploitation tradecraft should be deployed and tuned.

Default and built-in Fortinet accounts that are not operationally required should be disabled or removed. Any default account that is still required for operational reasons should be renamed to a non-predictable identifier and its access scope constrained to the minimum necessary for its operational function.

Strategic Considerations

FortiBleed is a symptom of what security researchers have documented as a broader challenge: configuration hygiene and credential lifecycle management that has not kept pace with the rate at which network security devices are deployed. The concentration of unrotated default credentials and legacy password hashing in a population of internet-facing enterprise firewalls indicates that deployment checklists are inconsistently applied and that ongoing credential maintenance is not systematically enforced across many organizations. Security teams should assess whether their organization maintains an authoritative inventory of all Fortinet devices, the firmware versions running on each, and the credential management status of administrative accounts. Where that inventory does not exist or is incomplete, establishing it is a prerequisite for effective response to incidents of this type.

Managed service providers and enterprises that operate FortiGate infrastructure on behalf of clients carry a heightened responsibility in the FortiBleed context. The credential dataset organized by target revenue suggests that threat actors have already prioritized follow-on campaigns against the most valuable environments. MSPs should proactively notify affected clients, provide evidence of credential rotation and MFA enforcement, and communicate any indicators of compromise found during their own log reviews.

The incidents documented in this note — multiple CVEs, a large-scale credential exposure, and a PBKDF2 migration that left legacy hashes accessible — provide grounds to revisit vendor security posture assessments. Vendor security posture, including the vendor's track record on timely patching, transparent disclosure, hardening by default, and responsiveness to discovered weaknesses, should be a factor in ongoing supplier risk assessments and contract renewals.

CSA Resource Alignment

The FortiBleed incident illustrates control failures at the intersection of identity management, network perimeter security, and vulnerability governance — each of which is addressed by CSA's published guidance and frameworks.

Zero Trust Architecture. CSA's Zero Trust guidance directly addresses the failure mode at the core of FortiBleed: the implicit trust extended to administrative credentials presented at a management interface. A Zero Trust approach removes the assumption that a device inside the network perimeter or authenticated through a VPN should receive elevated trust; instead, every administrative session is continuously authenticated, authorization is scoped to the minimum necessary, and privileged access is enforced through robust MFA. CSA's Zero Trust Advancement Center provides implementation guidance applicable to network perimeter device administration.

Cloud Controls Matrix (CCM). Multiple CCM control domains apply to this incident. IAM-02 (Strong Authentication) maps directly to the MFA enforcement gap. IAM-07 (User Access Reviews) covers the credential hygiene and account lifecycle failures that left default and inactive accounts in production. LOG-02 and LOG-09 address the log collection and review practices that would accelerate detection and response to unauthorized administrative sessions. CCM-mapped controls should be reviewed against FortiGate configuration baselines for all internet-facing deployments.

AI Controls Matrix (AICM). While FortiBleed is not an AI security incident in the conventional sense, the AI-assisted attack campaign documented by Amazon Threat Intelligence in parallel with FortiBleed is directly relevant to AICM's coverage of AI-enabled offensive operations. AICM addresses both the threat of AI-accelerated exploitation and the organizational controls needed to manage AI-powered adversarial activity in the threat landscape.

STAR Program. Organizations that use Fortinet infrastructure as part of cloud or hybrid environments that are STAR-registered should assess whether the FortiBleed exposure constitutes a material change to their security posture requiring disclosure or control re-attestation. The CSA STAR registry provides a mechanism for organizations to document their current security state and incident responses.

Shared Assessments and Supplier Risk. CSA's guidance on third-party risk and shared assessments is relevant to MSPs operating FortiGate on behalf of clients, as well as to enterprises that rely on Fortinet as a critical network security supplier. The FortiBleed timeline — encompassing multiple CVEs, a large-scale credential exposure, and a delayed PBKDF2 migration that left legacy hashes exploitable — warrants inclusion in supplier security reviews.

References

[1] SOCRadar. "FortiBleed 2026: 86,644 Fortinet Firewalls Compromised – Active Leak." SOCRadar, June 2026.

[2] Arctic Wolf. "Active FortiBleed Campaign Impacting Fortinet Devices Across 194 Countries." Arctic Wolf, June 2026.

[3] Help Net Security. "74,000 Fortinet firewall credentials exposed in FortiBleed data leak." Help Net Security, June 18, 2026.

[4] SOCRadar. "CVE-2026-24858: Patch Released for Fortinet FortiOS SSO Authentication Bypass." SOCRadar, January 2026.

[5] SecurityWeek. "Fortinet Patches Exploited FortiCloud SSO Authentication Bypass." SecurityWeek, January 2026.

[6] CISA. "CISA Urges Hardening Fortinet Devices After Reports of Credential Exposure." CISA Advisory, June 18, 2026.

[7] Kudelski Security. "Fortinet 'FortiBleed' Global Compromise & Active Exploitation of Fortinet Vulnerabilities." Kudelski Security Research Center, June 2026.

[8] The Hacker News. "AI-Assisted Threat Actor Compromises 600+ FortiGate Devices in 55 Countries." The Hacker News, February 2026.

[9] Fortinet Community. "Technical Tip: Enforcing PBKDF2 as hash function for administrator accounts in FortiOS v7.2.11 and later." Fortinet Community, 2025.

[10] Fortinet. "Fortinet 2026 Global Threat Landscape Report Reveals a Surge in AI-Enabled Cybercrime, Contributing to a 389% Increase in Ransomware Victims Year-over-Year." Fortinet Press Release, 2026.