CSAI Foundation | Cloud Security Alliance

FortiBleed: Mass VPN Credential Exposure at Enterprise Perimeters

2026-06-21

AI-assisted Rapid Research

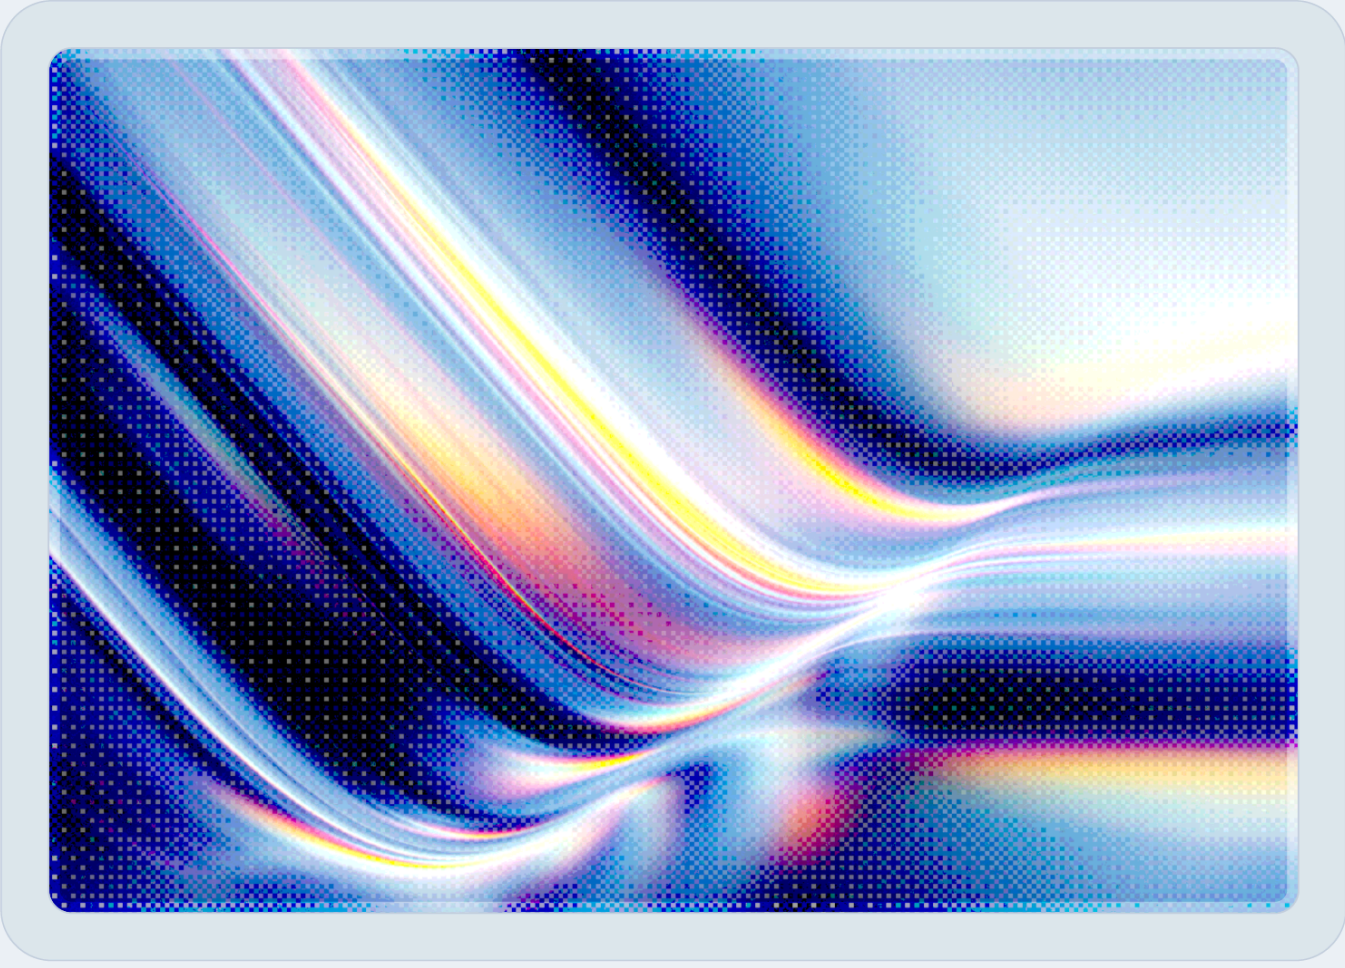

CSAI

CSA cloud security alliance®

Key Takeaways

- A credential harvesting campaign dubbed FortiBleed has produced a dataset of working administrative and SSL VPN credentials — sampled credentials confirmed authentic by independent researchers — for between 73,932 and 86,644 internet-facing Fortinet FortiGate appliances spanning 194 countries and 21,632 organizational domains [1][4][5].
- The dataset was discovered June 17, 2026, and is circulating in criminal underground communities; CISA issued an emergency alert on June 18, 2026 urging immediate session termination and credential rotation across all internet-facing Fortinet infrastructure [2], and Canada's Centre for Cyber Security issued a parallel alert the same day [16].
- FortiBleed exploits a convergence of conditions — credential stuffing using Fortinet-specific infostealer data, SSL VPN hash interception, and offline GPU cracking — compounded by a large proportion of devices still storing passwords as legacy SHA-256 hashes that modern cracking clusters can defeat in hours [8][12].
- Named organizations whose credentials appear in the dataset span Fortune 500 enterprises across energy, technology, telecommunications, manufacturing, and professional services, as well as at least one NATO-affiliated defense contractor [6][7].
- Organizations with internet-facing FortiGate appliances should immediately verify whether their devices appear in the exposed dataset using the lookup services described in the Recommendations section. Organizations confirmed as affected — or unable to rule out exposure — should treat this as an active incident and initiate the steps below. All organizations, regardless of dataset confirmation, should audit PBKDF2 hashing status and MFA coverage as a precautionary measure.

Background

Fortinet's FortiGate product line is one of the most widely deployed network security appliances in the world, providing unified threat management, firewall, intrusion prevention, and — most critically in this incident — SSL VPN remote access to hundreds of thousands of enterprise and government networks. This near-ubiquity in enterprise perimeter architecture makes FortiGate a high-value, high-yield target for mass credential harvesting: a single, well-curated credential list can enable access to a broad cross-section of internal networks across industries and geographies simultaneously.

On June 17, 2026, security researcher Volodymyr "Bob" Diachenko discovered a misconfigured attacker-controlled server hosting a large, structured dataset of what appeared to be active Fortinet FortiGate credentials [1][5]. The dataset was independently verified by threat intelligence firm Hudson Rock and researcher Kevin Beaumont, both of whom confirmed that sampled credentials were authentic and in active use [1][6]. The security community named the campaign FortiBleed, evoking both the Fortinet brand and the concept of credential leakage from a previously intact perimeter. CISA responded with an urgent alert on June 18, 2026, and Canada's Centre for Cyber Security issued Alert AL26-014 on the same date, both warning that the exposed credentials posed an immediate remote-access threat to affected organizations [2][3][16].

The FortiBleed disclosure arrived against a backdrop of pre-existing Fortinet security pressure. CVE-2026-24858, a FortiCloud SSO authentication bypass rated CVSS 9.8, had been identified as actively exploited and added to CISA's Known Exploited Vulnerabilities catalog on January 27, 2026 [9][10]. That vulnerability — which allowed any FortiCloud account holder to authenticate to devices registered to other users when FortiCloud SSO was enabled — broadened the attack surface available to threat actors during the period when the FortiBleed credential set was assembled. Researchers assess CVE-2026-24858 as a contributing, if not primary, factor in the underlying exposure [8][12].

Security Analysis

Attack Methodology

Analysis of artifacts recovered from the attacker-controlled server, reported by Diachenko and corroborated by multiple security firms, describes a methodical, multi-phase automated operation — evidence pointing to a deliberate, targeted campaign rather than opportunistic scanning [8][12]. The operators began by sweeping approximately 59.3 million internet hosts to enumerate exposed FortiGate management and VPN interfaces. Against the 320,777 FortiGate targets identified in this sweep, the attackers executed approximately 1.16 billion credential attempts [8]. Researchers assess this volume as consistent with pre-built, Fortinet-targeted credential lists sourced from prior breach data and infostealer logs, rather than generic password spray lists, based on the high success rate observed relative to attempt volume [8].

Where standard credential stuffing failed, the operators employed SSL VPN hash interception, capturing authentication hashes from session-level traffic and submitting them for offline cracking via a 45-GPU distributed cluster managed through Hashtopolis [8]. This technique bypasses online rate limiting entirely, since the hash cracking occurs on infrastructure the attacker controls rather than against the target device. Upon authenticating successfully, the operators moved laterally from the compromised

VPN endpoint into the connected Active Directory environment, escalating from perimeter access to internal network footholds [1][12]. Attribution based on language characteristics observed in attacker artifacts and forum communications points to a Russian-speaking multi-operator cybercriminal group, though attribution in credential harvesting campaigns of this type should be treated as provisional [1][7].

Technical Root Causes

FortiBleed's effectiveness rests on a combination of technical conditions, most prominently a legacy password hashing issue in FortiOS. Older FortiOS versions stored administrator passwords using SHA-256, a general-purpose cryptographic hash that is poorly suited to password storage because modern GPU clusters can evaluate billions of hash candidates per second [12][13]. The critical compounding factor is that the SHA-256 hash persists for existing accounts even after upgrading to a newer FortiOS version: the password is only re-hashed to the modern PBKDF2 algorithm when the administrator logs in following the upgrade [12]. An organization that upgraded FortiOS without subsequently triggering a credential re-hash for all administrator accounts remained exposed to GPU-accelerated offline cracking despite running a current software version. Fortinet addressed this by introducing PBKDF2 as the default hashing algorithm in FortiOS versions 7.2.11, 7.4.8, and 7.6.1, but the transition is not automatic — it requires administrative action to complete [12].

CVE-2026-24858 contributed a second, independent attack surface. The FortiCloud SSO authentication bypass — rated CVSS 9.8 and classified as CWE-288 (Authentication Bypass Using an Alternate Path or Channel) — allowed any FortiCloud account holder to authenticate to other users' registered devices when FortiCloud SSO was enabled [9][10]. This feature is activated automatically when an administrator registers a device to FortiCare via the GUI unless the administrator explicitly disables the "Allow administrative login using FortiCloud SSO" option, meaning many organizations had SSO enabled without deliberate intent. Fortinet patched the vulnerability in FortiOS 7.4.11 and released patches for additional affected products shortly thereafter [10][11]. Researchers estimated that the exposed dataset covers approximately 50% of all internet-reachable FortiGate devices [6], suggesting that a substantial proportion had not applied patches that would have mitigated the SHA-256 hashing exposure.

Scope and Impact

The FortiBleed dataset covers verified credentials for between 73,932 and 86,644 FortiGate devices — the range reflects differing verification methodologies across reporting firms — spanning 21,632 unique organizational domains in 194 countries [1][4][5][14]. Researchers estimate the affected population represents approximately 50% of all internet-reachable FortiGate devices worldwide, though this figure should be treated as an approximation pending more complete independent verification [6]. The named

organizations whose credentials appear in the dataset include Chevron and Sinopec in energy; Samsung, Foxconn, Lenovo, and Toyota in manufacturing and technology; Comcast and AT&T in telecommunications; Oracle and Accenture in technology services; and Siemens and PwC in industrial and professional services [6][7]. A Turkish NATO-affiliated defense contractor was also identified in the dataset [14], underscoring that the campaign reached critical infrastructure and defense sector targets alongside commercial enterprises.

The practical threat extends beyond the immediate credential exposure. Stolen but valid credentials enable VPN sessions that authenticate normally, meaning they may not trigger authentication anomaly alerts and may blend with legitimate remote-access activity. Lateral movement from a VPN endpoint into Active Directory — observed in this campaign — can give operators access to internal systems, additional credentials for further pivoting, and opportunities to establish persistence well removed from the initial entry vector. Organizations where VPN connections land on a segment with broad access to Active Directory face a significantly larger blast radius from a compromised credential, since lateral movement into internal systems requires no additional access boundary crossing.

Recommendations

Immediate Actions

Organizations with internet-facing FortiGate devices should first verify whether their systems appear in the exposed dataset using the lookup services published by multiple security vendors following the disclosure (see the Short-Term Mitigations section below). Organizations confirmed as affected — or unable to rule out exposure — should treat this as an active incident. All organizations, regardless of dataset confirmation, should audit PBKDF2 hashing status and MFA coverage as a precautionary measure. CISA's June 18, 2026 alert sets out the following recommended baseline steps [2]:

- Terminate all active SSL VPN and administrative sessions on all internet-facing Fortinet appliances.
- Reset all Fortinet VPN and administrator account passwords, with priority on systems accessible from the internet.
- Confirm that PBKDF2 is in use for administrator credential storage per Fortinet's guidance; if legacy SHA-256 hashes remain for any account, trigger re-hashing by having all administrators log in with their new credentials following the reset.
- Review authentication and session logs for suspicious activity: VPN sessions from unexpected geographic origins, sessions at anomalous hours, or sessions exhibiting rapid lateral

movement post-authentication.

- Enable phishing-resistant multi-factor authentication on all administrative and VPN access paths.

Short-Term Mitigations

Within the next two weeks, organizations should audit FortiOS patch levels across their entire Fortinet deployment. Devices not yet running FortiOS 7.2.11, 7.4.8, or 7.6.1 — or later maintenance releases — should be prioritized for upgrade to ensure PBKDF2 is operational as the default hashing scheme going forward [12]. CVE-2026-24858 remediation should be independently verified: any device that has not yet applied the relevant FortiOS patch, and any device where FortiCloud SSO remains enabled without a clear operational justification, represents an unresolved attack surface [9][10].

Network segmentation review is a high-priority parallel activity. Organizations where VPN connections land on a segment with broad access to Active Directory face a significantly larger blast radius from a compromised credential, since lateral movement into internal systems requires no additional access boundary crossing. Organizations should evaluate whether their VPN landing zones enforce least-privilege access, whether jump hosts or additional authentication gates separate VPN egress from internal critical systems, and whether endpoint detection tooling monitors for anomalous post-authentication behavior within the VPN-connected segment.

Organizations should additionally check their known FortiGate IP addresses and domain names against the breach notification and lookup services that multiple security vendors have published in response to FortiBleed, enabling direct confirmation of whether specific devices appear in the compromised dataset [15].

Strategic Considerations

FortiBleed is the latest in a recurring pattern of large-scale, targeted credential harvesting against SSL VPN infrastructure. Perimeter-exposed VPN appliances occupy a structurally dangerous position in enterprise architecture: they are internet-facing by necessity, they authenticate users for broad internal network access, and they are frequently patched on conservative maintenance cycles due to availability constraints, change management overhead, and the operational risk of appliance reboots in production environments. This incident provides a concrete case for treating network security appliances as software components subject to the same urgent patching standards as externally exposed application servers.

The role of infostealer malware in the FortiBleed attack chain deserves specific attention. The attackers did not rely on generic password lists; they used credentials known to have previously worked against Fortinet infrastructure, sourced from infostealer logs. When an endpoint infection affects a device with enterprise credentials stored — particularly credentials for perimeter devices, VPN appliances, or privileged administrative accounts — organizations should review whether those credentials were harvested and initiate targeted rotation. FortiBleed illustrates how infostealer logs, when assembled at scale, enable targeted credential-stuffing against specific infrastructure types; organizations should ensure their incident response procedures account for this downstream risk.

Over the longer term, organizations should evaluate their architectural dependence on traditional SSL VPN in light of Software Defined Perimeter and Zero Trust Network Access alternatives. SDP's authenticate-before-connect model removes the VPN appliance as a credential-stuffing target, since network resources are invisible to unauthenticated requests. This does not eliminate authentication attack surfaces entirely — the identity provider still presents a login endpoint — but it significantly reduces the number of exposed interfaces and moves authentication to a purpose-built, hardened system rather than a network appliance. Organizations adopting SDP should apply the same credential hygiene and MFA standards to their IdP that FortiBleed demonstrates are needed for VPN appliances. CSA's published SDP guidance provides implementation pathways for organizations that are ready to begin this transition.

CSA Resource Alignment

FortiBleed has direct implications for several foundational CSA publications and frameworks.

CSA's Software Defined Perimeter publications — including the SDP Architecture Guide v3 and the joint SDP and Zero Trust implementation guidance — describe the architectural alternative to SSL VPN that FortiBleed exploited. SDP's authenticate-before-connect model removes the VPN appliance as a credential-stuffing target by making network resources invisible to unauthenticated requests. This does not eliminate authentication attack surfaces entirely — the identity provider still presents a login interface — but it substantially narrows the exposed attack surface and places authentication responsibility on a purpose-built, hardened system rather than a network appliance. Organizations undertaking post-FortiBleed architecture reviews should consult these CSA publications as a foundation for perimeter modernization planning.

The AI Controls Matrix (AICM), CSA's cloud and AI security controls framework, contains IAM domain controls directly relevant to the FortiBleed threat vectors: multi-factor authentication requirements, privileged access management standards, credential lifecycle management, and network access control

minimization. Security teams should review their AICM control implementations against the specific gaps FortiBleed exploited — particularly around patching cadence for perimeter appliances and MFA coverage for VPN access paths.

The CSA STAR program provides a structured mechanism for organizations to publish and review security posture documentation against CSA-aligned frameworks. In the context of FortiBleed, organizations should assess whether their current STAR self-assessments accurately reflect the patch status and MFA coverage of their VPN and remote access infrastructure, and plan to update those assessments following remediation. Organizations evaluating vendor-managed Fortinet deployments should also request updated STAR documentation from their managed security service providers demonstrating that FortiBleed remediation is complete.

CSA's Zero Trust guidance — particularly its emphasis on micro-segmentation and least-privilege network access — addresses the blast radius component of this incident. Lateral movement from a compromised VPN endpoint is significantly more difficult in architectures that enforce Zero Trust principles and micro-segmentation within the internal network, since access to each resource requires separate authentication and authorization rather than relying on network adjacency.

References

[1] BleepingComputer. "FortiBleed leak exposes Fortinet VPN credentials for 73,000 devices." BleepingComputer, June 2026.

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

[3] BleepingComputer. "CISA warns Fortinet users to secure devices after FortiBleed leak." BleepingComputer, June 2026.

[4] SecurityWeek. "FortiBleed: 86,000 Fortinet Device Credentials Compromised." SecurityWeek, June 2026.

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

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

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

[8] BitSight. "Security Alert: Fortinet VPN Credentials Exposed." BitSight, June 2026.

[9] NIST National Vulnerability Database. "CVE-2026-24858 Detail." NVD, 2026.

[10] CISA. "Fortinet Releases Guidance to Address Ongoing Exploitation of Authentication Bypass Vulnerability CVE-2026-24858." CISA, January 28, 2026.

[11] Help Net Security. "Fortinet starts patching exploited FortiCloud SSO zero-day (CVE-2026-24858)." Help Net Security, January 28, 2026.

[12] Thrive. "The FortiBleed Autopsy — How 86,644 'Patched' Firewalls Became a Russian-Speaking Hacking Franchise." Thrive NextGen, June 2026.

[13] CybelAngel. "FortiBleed: 6 Things to Know About the Fortinet Credential Leak." CybelAngel, June 2026.

[14] Recorded Future. "FortiBleed Campaign Exposing Credentials for 73,932 FortiGate Systems." Recorded Future, June 2026.

[15] eSentire. "FortiBleed Campaign: Credentials Exposed for +80,000 Fortinet Appliances." eSentire, June 2026.

[16] Canadian Centre for Cyber Security. "AL26-014 – FortiBleed leak of thousands of compromised credentials impacting Fortinet devices." CCCS, June 18, 2026.