


FortiBleed: Credential Campaign and Admin Lockout at the Perimeter

Security Implications and Guidance

2026-10-09

 AI-assisted Rapid Research



CSAI

cloud
CSA security
alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

The FBI and U.S. Secret Service attribute FortiBleed primarily to credential reuse and weak password storage rather than to a specific software vulnerability [1]. They describe it as an active, global operation against internet-facing Fortinet FortiGate firewalls and SSL VPN gateways that relies on reused or leaked credentials and on legacy SHA-256 password storage [1]. The advisory cites SOCRadar's verification of more than 86,644 compromised devices across 194 countries [1][12]. Recorded Future reported a lower figure of 73,932 FortiGate systems with exposed credentials from its own analysis, so published totals differ by source, methodology, and the unit being counted [2].

The newest development is operational rather than technical. A joint FBI and Secret Service advisory dated October 6, 2026 reports that some victims have been locked out of their own devices because the operators create new administrator accounts and, in certain cases, delete existing accounts or change their passwords [1]. Organizations that assumed a patch or a routine password reset would close the exposure may therefore find they no longer control the device they are trying to remediate.

The campaign also feeds ransomware. According to the advisory, initial access brokers using the FortiBleed attack chain have passed access to ransomware affiliates, including INC/Lynx and Payload as of the advisory's publication [1]. A compromised firewall should therefore be treated as a potential precursor to domain-wide compromise, and not as a contained appliance problem.

Defenders should prioritize removing management interfaces from the internet, terminating all administrative and VPN sessions, resetting every credential, enforcing phishing-resistant MFA, and auditing device accounts and API keys for entries they did not create [1].

Background

FortiBleed became public in June 2026. Recorded Future's timeline places the first malicious HTTP activity on June 7, an underground auction of roughly 34,000 lines of VPN data on June 12, and public disclosure by researcher Volodymyr Diachenko on June 13, followed by validation from other researchers [2]. Fortinet's PSIRT published its own analysis of the reported credential compromise on June 19, 2026,

and stated that this is not a new Fortinet vulnerability [3]. The campaign's internal workflow became visible only because the operators unintentionally exposed a backend server as an open directory, which gave researchers an end-to-end view of tooling, datasets, and cracking infrastructure [1].

The U.S. government's October advisory speaks directly to persistence. Nearly four months after the initial disclosure, the FBI and Secret Service report that scanning of exposed Fortinet firewalls continues using previously obtained credentials [1]. The Record reported that CISA had issued an initial warning in June, and characterizes the October advisory as adding the lockout reporting that the earlier warning did not contain [4]. The Hacker News similarly reports that the campaign remains active and that FortiBleed-derived access continues to be monetized [5].

The name suggests a memory-disclosure flaw, but no CVE or patch is associated with the credential exposure itself [3][7]. Related vulnerabilities, such as the FortiCloud SSO authentication bypass tracked as CVE-2026-24858, have been discussed alongside the campaign in CSA's earlier analysis [6], although the federal advisory attributes the campaign's scale to credential reuse and weak hash storage and not to a specific software flaw [1]. This distinction matters for response: organizations that are fully patched can still be exposed.

Published totals also count different things, which is worth keeping in mind when comparing them. SOCRadar's figure refers to working device credentials and compromised devices as of June 19, 2026 [5][12]. Recorded Future's 73,932 counts FortiGate firewall URLs for which credentials were exposed [2], and its 320,777 figure counts targets of credential attempts [2]. CSA's earlier analysis of the 110-million figure counts credential records rather than devices [10].

Security Analysis

The Attack Chain

The FBI and Secret Service summarize the operators' progression in five broad stages based on artifacts left on the exposed server [1]. The operators scanned the internet for exposed FortiGate SSL VPN portals, then collected credentials through credential stuffing and password spraying built from earlier Fortinet leak dumps and infostealer logs. Password hashes obtained from compromised devices were funneled into a GPU-accelerated cluster, where Hashcat and Hashtopolis distributed the offline cracking work. Cracked credentials were enriched and validated, with scripts that filtered out honeypots and prioritized targets by revenue and network structure. The operation ended with packaged access sold to downstream actors, which is consistent with an initial access broker business model [1].

Public technical reporting adds detail on tooling and scale. Picus Security describes a Go-based sniffer, referred to as FortigateSniffer, that abuses the built-in FortiOS packet-capture diagnostic to passively collect authentication traffic across 24 protocols, including NTLM, Kerberos, and RADIUS [7]. Recorded Future reports 1.16 billion credential attempts against 320,777 FortiGate targets and a 45-GPU cluster managed through Hashtopolis [2]. These figures come from vendor analysis of the recovered artifacts, and readers should treat them as reported rather than independently audited.

Why Legacy Hashing Prolongs Exposure

A central enabling factor is how older FortiOS releases store administrator passwords. The FBI and Secret Service cite legacy SHA-256 storage as one reason cracked credentials were obtainable at scale, and they direct organizations to enforce PBKDF2 for administrator accounts and remove weaker legacy hashes [1]. Fortinet's guidance covers FortiOS 7.2.11 and later [8]. Picus notes a caveat that is easy to miss: PBKDF2 takes effect only after each administrator re-authenticates, so a device can be upgraded and still hold crack-friendly hashes for accounts that have not logged in since [7]. This suggests that "we upgraded" is not sufficient evidence of remediation, and that credential storage should be verified per account.

Persistence and the Lockout Mechanism

The lockout behavior builds on the persistence technique. During intrusion, the operators create administrator accounts that did not previously exist on the device, and the advisory lists account names observed on victim systems, among them adminin, fortiAdmin, forticloud-sync, fgtsecure, and support_fortinet [1]. Several of these names imitate vendor or support functions; in our assessment, this naming is likely intended to survive casual review. In certain cases, the actors delete existing accounts or change their passwords, mapped by the advisory to ATT&CK technique T1531 (Account Access Removal), to block the owner from the device while lateral movement continues [1].

The practical consequence is that standard remediation sequencing can fail. If a team plans to reset passwords from the device's management plane, and the legitimate accounts have already been removed, recovery may require console or out-of-band access, restoration from a known-good configuration, or vendor support. The advisory also warns that REST API keys, which enable automated configuration and backup, can serve as a persistence path, and it recommends removing unknown keys and refreshing legitimate ones [1]. Organizations that rotate human credentials alone may leave a working access path in place.

Downstream Impact

The perimeter device matters to attackers because of what it sees and where it sits. With verified credentials, the operators enumerated Active Directory and sprayed passwords to find privileged accounts [1]. The sniffer's collection of NTLM, Kerberos, and RADIUS traffic, as described by Picus, extends the harvest beyond the firewall's own accounts to credentials belonging to internal users [7]. The Record reports at least a dozen organizations that were breached and encrypted [4]. It also relays an allegation that Russian actors used FortiBleed access to reach U.K. government email accounts; the source describes this as unconfirmed, and it should be treated as unverified here [4].

Summary of Reported Indicators and Controls

The following table maps the main elements of the campaign to the observable evidence and the corresponding control. It draws on the federal advisory [1] and on our own analysis, and some rows, such as the out-of-band access plan and VPN pool segmentation, are our synthesis rather than advisory text.

Campaign element	Observable evidence	Primary control
Internet-exposed management and SSL VPN portals	Repeated scanning and brute-force attempts in firewall and VPN logs	Restrict management via trusted hosts, local-in policy, or remove internet administration
Credential stuffing and spraying	Successful logins from unfamiliar IP addresses	Phishing-resistant MFA on all remote and administrative accounts
Offline hash cracking	Legacy SHA-256 administrator hashes	Enforce PBKDF2; require administrators to re-authenticate
Persistence	New administrator accounts, unfamiliar REST API keys	Compare configuration to a known-good baseline; remove unknown keys
Lockout	Original accounts deleted or passwords changed	Out-of-band access plan; configuration restore procedure
Lateral movement and sale of access	AD enumeration, spraying from device-originated paths	Review domain controller and authentication logs; segment VPN pools

The advisory publishes IP addresses and account names for hunting, and it cautions that infrastructure may be dynamically reassigned and that indicators should be vetted against current telemetry before blocking [1].

Recommendations

Immediate Actions

Remove FortiGate administrative interfaces from direct internet exposure. The advisory ranks the options as trusted hosts (good), a local-in policy (better), and no internet administration at all (best) [1]. Terminate all active administrative and SSL VPN sessions, then reset all VPN and administrative passwords, starting with internet-facing systems [1]. Require phishing-resistant MFA on every remote access and administrative account [1].

Before relying on any reset, confirm that you still control the device and that the account list is legitimate. Review all local accounts against the published names and against your own records, review configuration changes against a known-good baseline, and inventory and refresh REST API keys [1]. If original administrator accounts are missing or no longer accept known credentials, treat the device as compromised and plan recovery through console access or a rebuilt configuration. Review firewall, VPN, authentication, and domain controller logs for lateral movement, and report relevant findings to the FBI's IC3 or a local field office [1].

Short-Term Mitigations

Verify credential storage per administrator account and force re-authentication where PBKDF2 has not yet taken effect [7][8]. Rotate secrets that may have crossed the device, including service accounts, LDAP bind credentials, RADIUS shared secrets, and any directory credentials that authenticated through it, because the sniffer reportedly captured internal authentication protocols [7]. Reset the passwords of users who authenticated through exposed VPN portals, and check infostealer-exposure data for those users, since leaked logs seeded the stuffing attacks [1]. Pre-stage a documented lockout-recovery runbook that includes out-of-band console access and offline copies of known-good configurations.

Strategic Considerations

FortiBleed illustrates that a patched appliance can still be an exposed one when authentication relies on reusable secrets and the management plane is reachable from the internet [1][7]. Organizations should treat perimeter devices as identity systems, with the same lifecycle controls applied to administrator credentials, API keys, and break-glass accounts that they apply to cloud consoles. Architectures that authenticate users before exposing network paths, such as software-defined perimeter and zero trust network access models, can reduce the value of a single stolen VPN credential. In our assessment, concentration risk also deserves attention: when a large share of an organization's remote access depends on one vendor's appliance family, a campaign of this kind could affect all of it simultaneously.

CSA Resource Alignment

CSA has published several analyses on this campaign during the summer. The closest match to this note is *FortiBleed: Mass VPN Credential Exposure at Enterprise Perimeters* [6], which covers the initial credential dataset, the hashing weakness, and recommendations for patching, MFA, and segmentation. This note updates that analysis with the October federal advisory's lockout and persistence findings, which arrived after the earlier work was published. Organizations that followed the earlier guidance on password resets should revisit it in light of the account-deletion behavior described above.

FortiBleed: Weaponized FortiGate Firewalls in Mass Credential Harvest [9] examines how compromised firewalls themselves become collection platforms, which aligns with the sniffer and internal-protocol harvesting described here. *FortiBleed: Anatomy of a 110-Million-Credential Harvesting Campaign* [10] provides additional scale and campaign-structure context that complements the stage-by-stage view in the federal advisory. *FortiBleed: Default Credential Exploitation and Mass Fortinet Compromise* [13] is also relevant to the credential-reuse thesis.

For control mapping, the AI Controls Matrix v1.1 [11] is a superset of the Cloud Controls Matrix and provides a vendor-neutral structure for the identity and access management, threat and vulnerability management, and logging domains touched by these recommendations. Organizations redesigning remote access should apply it together with CSA's Zero Trust and software-defined perimeter guidance.

References

- [1] FBI and U.S. Secret Service. "[FortiBleed Operations Continue Targeting Exposed Systems Leading to Reports of Lockouts \(JCSA-20261006-01\)](#)." Internet Crime Complaint Center, October 6, 2026.
- [2] Recorded Future. "[FortiBleed Campaign Exposing Credentials for 73,932 FortiGate Systems](#)." Recorded Future, 2026.
- [3] Fortinet PSIRT. "[Analysis of Reported Credential Compromise of FortiGate Devices](#)." Fortinet, June 19, 2026.
- [4] The Record. "[FBI, Secret Service add to warnings of FortiBleed credential stealing campaign](#)." Recorded Future News, October 2026.
- [5] The Hacker News. "[FBI Warns FortiBleed Remains Active After Amassing 86,644 Fortinet Device Credentials](#)." The Hacker News, October 2026.
- [6] Cloud Security Alliance. "[FortiBleed: Mass VPN Credential Exposure at Enterprise Perimeters](#)." CSA Labs, June 2026.
- [7] Picus Security. "[FortiBleed: Inside the Campaign That Cracked 75,000 Fortinet Firewalls](#)." Picus Security, 2026.
- [8] Fortinet. "[Technical Tip: Enforcing PBKDF2 as hash function for administrator accounts in FortiOS v7.2.11 and later](#)." Fortinet Community, 2026.
- [9] Cloud Security Alliance. "[FortiBleed: Weaponized FortiGate Firewalls in Mass Credential Harvest](#)." CSA Labs, June 24, 2026.
- [10] Cloud Security Alliance. "[FortiBleed: Anatomy of a 110-Million-Credential Harvesting Campaign](#)." CSA Labs, June 25, 2026.
- [11] Cloud Security Alliance. "[AI Controls Matrix \(AICM\) v1.1](#)." Cloud Security Alliance, June 22, 2026.
- [12] SOCRadar. "[FortiBleed: Fortinet Firewalls Compromised](#)." SOCRadar, 2026.
- [13] Cloud Security Alliance. "[FortiBleed: Default Credential Exploitation and Mass Fortinet Compromise](#)." CSA Labs, June 2026.