CSAI Foundation | Cloud Security Alliance

FortiBleed: Weaponized FortiGate Firewalls in Mass Credential Harvest

Russian IAB deploys FortigateSniffer across 430,000+ devices, harvesting over 110 million credentials across 24 protocols

2026-06-24

AI-assisted Rapid Research

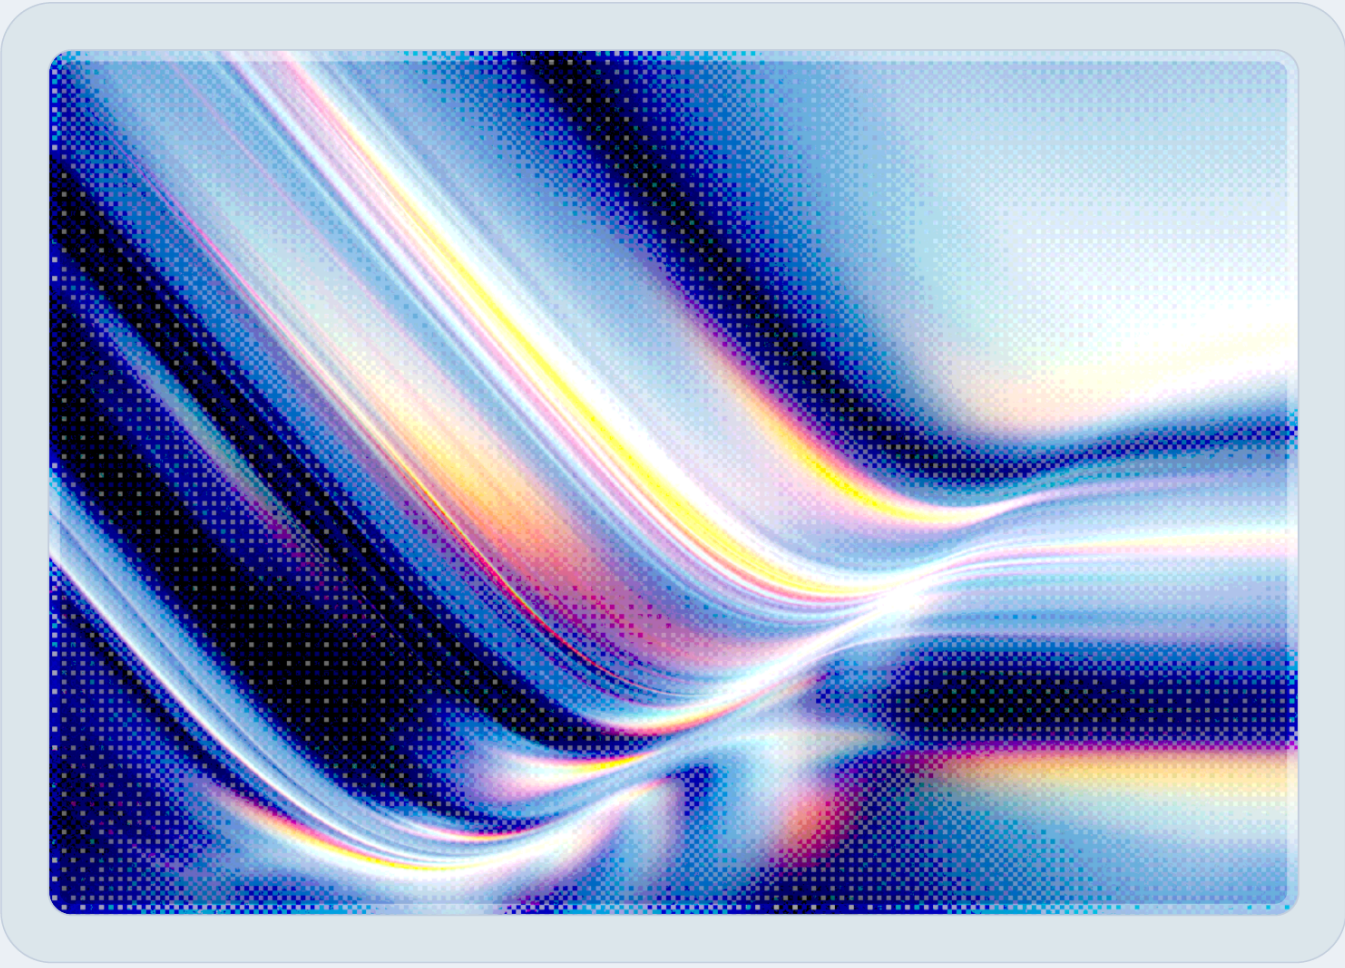

CSAI

CSA cloud security alliance®

Key Takeaways

- A likely Russian-speaking, financially motivated initial access broker — identified through Russian-language tooling and Moscow-Time operational scheduling [4][5] — has operated the FortiBleed campaign since at least February 2026, compromising more than 430,000 FortiGate firewalls globally and harvesting over 110 million credentials across 24 authentication protocols, subsequently monetizing access through underground forums. [1][2][3]
- The campaign's core implant, FortigateSniffer, is a custom Golang-based tool that abuses the legitimate FortiOS built-in diagnostic command `diagnose sniffer packet` to passively intercept authentication traffic from compromised devices — without deploying traditional malware or triggering endpoint defenses. [4][5]
- Initial access relied primarily on CVE-2026-24858, a critical (CVSS 9.4) FortiCloud SSO authentication bypass enabling rogue administrative account creation on fully patched devices, combined with credential stuffing against the approximately 63% of firewalls running default or generic administrator credentials. [6][7]
- Roughly 50% of all internet-reachable FortiGate devices are assessed to be affected across 194 countries; small and medium-sized businesses with fewer than 200 employees represent the predominant target pool, with confirmed follow-on impact reaching a NATO-aligned defense contractor. [2][8]
- Researchers found no evidence of zero-day vulnerability exploitation in this campaign; initial access relied on a patched authentication bypass (CVE-2026-24858) and credential stuffing against default accounts. [6][7] Researchers attribute the campaign's broad success to persistent configuration weaknesses, poor credential hygiene, and legacy SHA-256 password hashes surviving FortiOS firmware upgrade paths. [7]
- The same infrastructure and tradecraft have been simultaneously directed against Synology NAS devices, Sophos firewalls, RDWeb portals, Citrix SSL-VPN gateways, and MS-SQL servers since late February 2026, making this a multi-vendor perimeter-device threat. [2][3]

Background

FortiGate firewalls occupy a structurally privileged position in enterprise network architecture. Deployed as perimeter enforcement points, SSL-VPN concentrators, and network segmentation boundaries, they authenticate every remote user and proxy authentication traffic for downstream services. That structural role makes them ideal targets for a financially motivated threat actor: a single compromised device yields credentials for every user authenticating through it. Internet scan data consistently places FortiGate among the most widely internet-exposed network appliances globally, providing a large attack surface accessible to opportunistic and targeted actors alike. [2]

FortiBleed emerged against the backdrop of a well-established pattern of Russian-speaking criminal and state-adjacent actors targeting firewall and VPN appliances. The campaign is notable for both its operational scale — 659 simultaneous credential-harvesting pipelines and more than 430,000 compromised devices — and its reliance on repurposed diagnostic functionality to avoid detection. Between May 31 and June 15, 2026 alone, researchers documented 659 distinct credential-harvesting pipelines operating simultaneously. [3][4]

The campaign's origins trace to late January 2026, when Fortinet disclosed CVE-2026-24858, a critical authentication bypass in FortiCloud's Single Sign-On implementation. The vulnerability allowed any attacker holding a FortiCloud account to authenticate against devices registered to other FortiCloud tenants, enabling administrative account creation on fully patched systems without any credential knowledge. CISA added CVE-2026-24858 to its Known Exploited Vulnerabilities catalog on January 27, 2026, and directed Federal Civilian Executive Branch agencies to remediate within three days. [6] Despite CISA's January 27 remediation directive — one of the agency's shortest remediation windows — the threat actor had already established footholds across thousands of devices before the advisory was published. [6] The sniffer deployment phase accelerated through the following months.

Security Analysis

Initial Compromise: CVE-2026-24858 and Credential Weaknesses

The threat actor applied a layered initial access strategy. On devices with FortiCloud SSO enabled, CVE-2026-24858 provided a low-barrier pathway for any attacker with a FortiCloud account: they could authenticate against any other registered device, create a persistent local administrator account, and

export configuration files without triggering failed-authentication telemetry. Fortinet temporarily disabled the FortiCloud SSO service globally on January 26, 2026 to contain active exploitation before reinstating it with mitigations the following day. [9][6]

Where SSO exploitation was unavailable, the actor used a custom tool called "forticheck" for automated credential stuffing against management interfaces, combined with reconnaissance via Masscan and Shodan to enumerate internet-exposed devices. The credential stuffing achieved high success rates: researchers found that approximately 63% of the credentials recovered from compromised devices were either unchanged default Fortinet system accounts or generic administrative accounts using simple passwords, eliminating the need for brute force on the majority of targets. [7] This finding makes clear that configuration hygiene — not only vulnerability patching — constitutes a primary attack surface on network devices.

The FortigateSniffer Implant

Once administrative access was established, the threat actor deployed FortigateSniffer, a custom Golang-based tool compiled for both Linux (fg_sniffer_linux_amd64) and Windows (fg_sniffer_windows_amd64.exe) execution environments. The tool's user interface and embedded source code comments are entirely in Russian, providing linguistic indicators consistent with Russian-origin development. [4][5]

FortigateSniffer operates through a technique that avoids writing traditional malware to disk. It establishes an SSH connection to the compromised device and continuously invokes the FortiOS built-in command `diagnose sniffer packet` — a legitimate administrative utility intended for network troubleshooting. By invoking this command in an ongoing loop, FortigateSniffer passively captures all authentication traffic traversing the firewall across 24 protocols simultaneously. The monitored protocol list includes Kerberos, NTLM, LDAP, RADIUS, TACACS+, RDP, WinRM, SMB, RPC, FTP, Telnet, SMTP, IMAP, POP3, and multiple database protocols including MS-SQL, MySQL, and PostgreSQL. [5][3] Because the tool exercises a built-in FortiOS diagnostic facility rather than injecting custom code into the operating system, it leaves a minimal on-device artifact footprint and is unlikely to trigger signature-based endpoint detection designed for traditional malware. [4][5]

Captured traffic is processed by a companion component called SNIFTRAN, which reconstructs network flows into PCAP and PCAPNG format. These files are analyzed by a Python-based toolkit that extracts cleartext credentials, NTLM v2 hashes, Kerberos ticket material, session cookies, and database authentication tokens. The toolkit produces Hashcat-compatible output files partitioned by credential type for downstream cracking operations. [4]

Credential Processing Pipeline

Offline hash cracking employs a distributed GPU cluster managed through Hashtopolis, with Hashcat as the underlying engine; public reporting describes a substantial GPU fleet, though the precise count varies across analyses. [3] A Telegram bot designated HASHBOT provides the operator with live telemetry on cracking progress and pipeline status. An additional binary called CyberStrike Harvester v1.5 manages credential validation cycles, testing cracked credentials against live services at 1,000 simultaneous threads per cycle on a 300-minute interval schedule; the initial cycle success rate was reported at approximately 90%. [11]

The credential harvest totals over 110 million items in aggregate. Researchers have broken this figure down into approximately 14.8 million RADIUS credentials, 924,000 NTLM hashes, 130,000 Kerberos hashes, and 89 million MySQL authentication tokens, alongside additional material from other monitored protocols. [3] Access is monetized through underground forums; accounts were listed at an initial asking price of $30,000, which reportedly doubled to $60,000 within hours of publication. [4]

Analysts also identified evidence of deliberate backdoor provisioning: identical username-and-password pairs — such as "adminin:ITAdmin@888" — appearing identically across 3,947 distinct compromised devices, strongly suggesting the actor installed persistent access accounts intended to survive credential rotation if administrators failed to audit local account lists during remediation. [4]

Operational Security Tradecraft

The campaign demonstrated systematic operational security. Sniffer activity is constrained to the window of 7:00 a.m. through 6:00 p.m. Moscow Time, blending harvested traffic patterns with normal business-hours authentication volumes and complicating volume-based anomaly detection. The precision of this time-bound constraint — eleven hours aligned exactly to Moscow business hours — is a stronger attribution indicator than language artifacts alone, since timing parameters embedded in operational tooling are substantially harder to fabricate as false flags. Geofencing filters restrict collection to specific IP ranges, limiting exposure in unmonitored segments and focusing resources on economically ranked targets. [3][5] The presence of a single hardcoded Telegram administrator account in the HASHBOT configuration suggests individual operator control, though the supporting infrastructure — hundreds of cracking servers and 659 simultaneous pipelines — implies access to substantial shared criminal resources. [4]

Scale, Sectoral Impact, and Multi-Vendor Scope

FortiBleed's confirmed reach spans 194 countries. Primary target concentration is in the United States and India, with a pronounced focus on organizations with fewer than 200 employees — consistent with IAB monetization strategy, which favors volume access to smaller targets that can be bundled and resold to ransomware affiliates or intelligence customers. [2][3] The most significant confirmed victim at the time of this writing is a NATO-aligned defense contractor breached on June 15, 2026, from which DFS backup data was exfiltrated. This incident raises the possibility of intelligence-motivated tasking or opportunistic resale of sensitive access alongside the campaign's primary financial profile. [2]

The same actor's infrastructure has simultaneously targeted Synology NAS systems, Sophos firewalls, RDWeb portals, Citrix SSL-VPN gateways, and MS-SQL servers since at least February 28, 2026. Organizations that have resolved their FortiGate exposure but depend on these other perimeter or remote-access technologies should apply the same investigation and remediation posture described below. [2][3]

The Legacy Hashing Problem

A structural amplifier of the campaign's effectiveness is Fortinet's firmware upgrade path for password hashing. Fortinet introduced PBKDF2-based hashing — substantially more resistant to GPU-accelerated cracking than the previously used salted SHA-256 scheme — in FortiOS versions 7.2.11, 7.4.8, and 7.6.1. However, the PBKDF2 upgrade applies only to credentials that are actively re-authenticated following the firmware update. Accounts that have not re-authenticated after the upgrade retain SHA-256 hashes in configuration backups and persistent storage, and those hashes remain vulnerable to GPU-accelerated cracking, particularly where underlying passwords are weak or follow common patterns. [7] An organization that believes it has addressed credential security through firmware updates may still expose legacy service account or dormant administrator hashes if post-upgrade login was not enforced across all accounts.

Recommendations

Immediate Actions

Organizations running FortiGate devices should treat FortiBleed as an active incident requiring investigation, not only a patch-and-move-on advisory. Terminating all active SSL-VPN and administrative sessions immediately is the highest-priority first step, as it interrupts any FortigateSniffer

connections already established on devices. Every administrative password must be rotated using strong, unique credentials — not changed to a new generic value. Administrators should log back into devices post-firmware-upgrade to force re-hashing under PBKDF2, and the FortiOS `login-lockout-upon-weaker-encryption` setting should be enabled on all devices running v7.2.x or v7.4.x. [7]

All local account lists require a thorough audit. Any account not provisioned through the organization's documented processes — particularly accounts with names or passwords matching the patterns identified across thousands of compromised devices — should be treated as a planted backdoor and removed immediately. FortiCloud SSO should be disabled on any device not requiring it, and management interfaces should be removed from internet exposure wherever operationally feasible. [7]

Short-Term Mitigations

Enrolling all administrative and SSL-VPN accounts in phishing-resistant MFA — specifically FIDO2 or certificate-based authentication — is among the highest-leverage structural mitigations for credential-based intrusions of this type. The campaign's reliance on credential stuffing as the primary access vector for the majority of targets suggests that phishing-resistant MFA would have significantly reduced successful intrusions on devices where credentials were known or guessable; MFA alone does not address authentication bypass vulnerabilities such as CVE-2026-24858, which requires patching and SSO configuration controls. [7][9] Network access to firewall management interfaces should be restricted to dedicated administrative VLAN segments with no direct internet routing, ensuring that management traffic cannot originate from the same networks authenticated VPN clients reach. [7][9]

Log review should span SSL-VPN access logs, Active Directory authentication events, SMB share access records, and database authentication logs. Specific patterns meriting investigation include VPN session logins from unusual geographies, lateral movement activity originating from VPN-assigned IP ranges, and bulk recursive directory listing activity on file servers indicative of data staging. For organizations also running Synology NAS, Sophos firewalls, RDWeb, Citrix SSL-VPN, or MS-SQL infrastructure, the same credential rotation and MFA enforcement should be applied. [2][4]

Strategic Considerations

FortiBleed surfaces a structural tension in enterprise network architecture: perimeter security devices often receive implicit trust based on their role, yet may be among the least monitored systems in smaller organizations — and, as this campaign demonstrates, are high-value targets that warrant the same monitoring rigor applied to internal identity and endpoint infrastructure. A firewall that sits in the path of all authentication traffic is, by design, a credential interceptor if compromised. The threat actor exploited

this by repurposing built-in diagnostic functionality rather than deploying detectable malware — an approach that requires no novel exploitation and leaves no signature for conventional endpoint defenses to catch.

Organizations should establish and maintain authoritative inventories of all network appliances, tracking firmware versions and the date of last administrative credential rotation. Network devices should be integrated into SIEM and UEBA monitoring so that anomalous management-plane activity — unexpected outbound SSH connections, unusual configuration export events, elevated authentication traffic volumes, or diagnostic command invocations outside of change windows — triggers alerts. Managed service providers administering FortiGate fleets on behalf of clients bear heightened responsibility here: their own management plane credentials are a high-value target, and proactive client notification and evidence sharing with affected organizations are appropriate responses to a campaign of this scope.

CSA Resource Alignment

FortiBleed maps to several dimensions of CSA's published guidance. CSA's Zero Trust framework — which rejects implicit trust based on network location and treats all devices, users, and sessions as potentially compromised until verified — applies directly to the conditions the campaign exploited. Treating management-plane access to network appliances as inherently untrusted, and requiring verified identity and MFA before permitting configuration or diagnostic operations, represents the Zero Trust posture that would have disrupted this campaign's delivery mechanism. CSA's Software-Defined Perimeter work provides specific architecture guidance for removing management interfaces from public-internet exposure, which remains among the most effective structural controls for perimeter-device security.

The AI Controls Matrix (AICM) and Cloud Controls Matrix (CCM) contain directly applicable control domains covering privileged access management, credential lifecycle management, multi-factor authentication enforcement, and session management. Organizations using the STAR registry to communicate their security posture should treat network device credential management as a distinct audit scope. FortiBleed's finding that 63% of compromised credentials were default or generic accounts suggests this control domain is systematically under-addressed in standard organizational assessments and warrants explicit attention. [10]

CSA's research on shadow access and non-human identity provides useful framing for the planted-credential finding. Uniform credential pairs provisioned across thousands of devices represent a class of non-human, automated access that bypasses normal human identity lifecycle processes. Identity

governance workflows focused on human account provisioning and deprovisioning may not surface implanted device-level accounts without device-specific audit scope; organizations need account audit processes that operate at the network appliance layer.

References

[1] BleepingComputer. "FortiBleed campaign used custom FortiGate sniffer to steal credentials." BleepingComputer, June 2026.

[2] SecurityWeek. "Russian Initial Access Broker Behind FortiBleed Campaign." SecurityWeek, June 2026.

[3] The Hacker News. "FortiBleed Targeted FortiGate Firewalls in 110 Million-Credential Harvesting Operation." The Hacker News, June 2026.

[4] Security Affairs. "FortiBleed: The Most Detailed Breakdown Yet of an Active Russian Credential-Harvesting Operation." Security Affairs, June 2026.

[5] CyberSecurityNews. "Hackers Using FortigateSniffer Tool That Turns Compromised Firewalls Into Password Collectors." CyberSecurityNews, June 2026.

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

[7] Cloud Security Alliance Labs. "FortiBleed: Default Credential Exploitation and Mass Fortinet Compromise." CSA Labs, June 2026.

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

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

[10] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2025.

[11] Arctic Wolf. "Inside FortiBleed: Reverse Engineering the CyberStrike Harvester Behind a Global FortiGate Credential Factory." Arctic Wolf, June 2026.

