CSAI Foundation | Cloud Security Alliance

CVE-2026-0300: Root RCE Actively Exploited in PAN-OS

State-Sponsored Campaign Targets Firewall Authentication Portals via Buffer Overflow

2026-05-13

AI-assisted Rapid Research

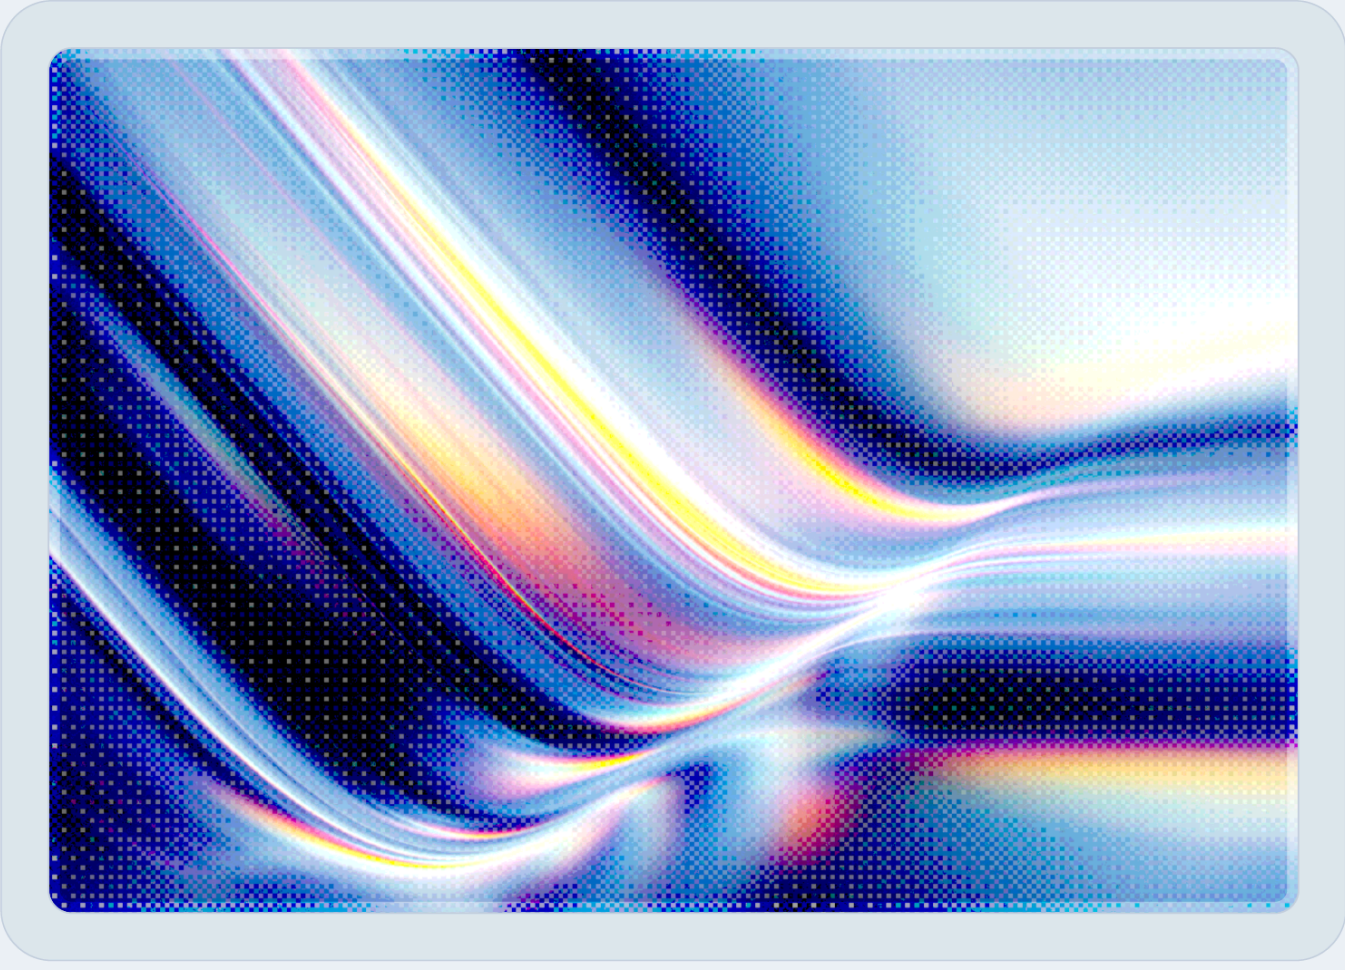

CSAI

CSA cloud security alliance®

Key Takeaways

- CVE-2026-0300 is an unauthenticated buffer overflow in the PAN-OS User-ID Authentication Portal carrying a CVSS 4.0 score of 9.3 (Critical), enabling root-level remote code execution on PA-Series and VM-Series firewalls without credentials or user interaction [1].
- Active exploitation began April 9, 2026—nearly a month before public disclosure—with confirmed successful compromise achieved on April 16; Unit 42 tracks the campaign as CL-STA-1132, a likely state-sponsored threat cluster with tactical overlaps to Volt Typhoon and APT41 [4].
- Post-exploitation activity includes deliberate forensic destruction, Active Directory enumeration, and deployment of open-source tunneling tools (EarthWorm, ReverseSocks5), indicating a persistent intelligence-collection mission rather than opportunistic access [4].
- CISA added CVE-2026-0300 to its Known Exploited Vulnerabilities catalog on May 6, 2026, requiring Federal Civilian Executive Branch agencies to apply mitigations by May 9, 2026 and complete full patching by May 27, 2026; software fixes began releasing May 13, 2026 [1][2].
- Organizations must immediately restrict or disable internet-facing User-ID Authentication Portals and conduct forensic investigation for pre-disclosure compromise, accounting for the attackers' documented log-clearing behavior [1][4].

Background

Palo Alto Networks PAN-OS is the operating system running PA-Series hardware firewalls and VM-Series virtual firewalls, deployed at network edges by more than 70,000 customers globally, including 90 percent of Fortune 10 companies and major financial institutions [3]. The platform's User-ID feature enables identity-aware policy enforcement by mapping authenticated user identities to IP addresses, and its Authentication Portal—also referred to as Captive Portal—plays a central role in that mapping. The portal intercepts connections from unauthenticated clients, presents a login interface, and passes confirmed identity information to the policy engine once authentication succeeds. Because the portal is designed to receive connections from clients who have not yet proven their identity, it must be exposed to the network segments containing those clients, creating an inherent tension between functional necessity and security boundary protection.

On May 6, 2026, Palo Alto Networks publicly disclosed CVE-2026-0300, a critical buffer overflow in the User-ID Authentication Portal component of PAN-OS [1]. The flaw enables an unauthenticated attacker to execute arbitrary code at the root privilege level on affected devices by sending specially crafted network packets. With a CVSS 4.0 score of 9.3—derived from a vector of `AV:N/AC:L/AT:N/PR:N/UI:N/VC:H/VI:H/VA:H/SC:L/SI:L/SA:N`—this vulnerability occupies the highest tier of severity in the CVSS scoring framework [1]. Network accessibility, low attack complexity, no authentication requirement, and high impact across confidentiality, integrity, and availability collectively produce the highest tier of CVSS scoring for network-accessible vulnerabilities. Disclosure followed Palo Alto Networks' internal detection of in-the-wild exploitation, and CISA's same-day addition of the vulnerability to its Known Exploited Vulnerabilities catalog confirmed the threat as active and immediate [2].

Affected products include all PA-Series hardware firewalls and VM-Series virtual firewalls running PAN-OS versions 10.2, 11.1, 11.2, or 12.1 with the User-ID Authentication Portal configured and network-accessible [1]. Prisma Access, Cloud NGFW, and Panorama management appliances are not affected. The exposure scope is substantial: Shadowserver data cited by BleepingComputer identified more than 5,800 VM-Series firewall instances currently reachable from the internet, concentrated in Asia (approximately 2,466 instances) and North America (approximately 1,998 instances), with PA-Series hardware appliances in similar configurations extending the attack surface further [3].

Security Analysis

Vulnerability Mechanics

CVE-2026-0300 is an out-of-bounds write vulnerability in PAN-OS's User-ID Authentication Portal, which operates as an nginx-hosted web service on the firewall [1]. When the portal processes certain specially crafted HTTP requests from unauthenticated clients, a flaw in memory boundary validation allows attacker-supplied data to be written beyond the bounds of an allocated buffer, corrupting adjacent memory regions. In the exploitation chain observed in the wild, attackers leveraged this memory corruption to inject shellcode directly into a running nginx worker process [4]. Because the portal service operates with root-level operating system privileges, the injected shellcode inherits those permissions, granting the attacker root-level control over the underlying firewall system—including its network interfaces, policy tables, and stored credentials.

The CVSS 4.0 scoring reflects the threat's practical characteristics accurately. The score drops from 9.3 to 8.7 when portal access is restricted exclusively to trusted internal IP addresses—evidence that network-level access controls are highly effective as an interim measure and align with sound long-term

architectural practice [1]. The vulnerability requires no special conditions, no existing user session, and no assistance from any person on the target network; an attacker with TCP connectivity to the portal's listening interface on any affected version can attempt exploitation directly.

Observed Exploitation Campaign

Unit 42 reconstructed a detailed timeline of the observed exploitation campaign, designating the activity cluster CL-STA-1132 [4]. The earliest recorded exploitation attempts occurred on April 9, 2026—approximately four weeks before public disclosure—with initial probes failing to produce a successful compromise. One week later, on April 16, the attackers achieved remote code execution, injecting shellcode into the nginx worker process and establishing a root-privilege foothold. By April 20, additional tooling had been deployed with elevated privileges, and the adversary began Active Directory enumeration, querying domain root structures and DomainDnsZones to map the victim organization's identity infrastructure. On April 29 the campaign escalated: a SAML flood attack was launched, a second PAN-OS device was compromised, and tunneling utilities were downloaded to establish persistent, covert communication channels [4].

The post-exploitation behavior throughout this campaign is consistent with a deliberate focus on operational persistence and forensic evasion. Immediately following successful compromise, the attackers engaged in systematic destruction of forensic artifacts: crash kernel messages were cleared, nginx crash logs and records were deleted, and core dump files that might have revealed exploitation details to incident responders were removed [4]. This evidence destruction, combined with the use of intermittent interactive sessions spread across multiple weeks and a deliberate preference for open-source tooling over custom implants, reflects an adversary whose primary operational concern is sustained, undetected access rather than rapid data exfiltration.

Threat Actor Profile and Tooling

Unit 42 characterizes CL-STA-1132 as a likely state-sponsored threat cluster based on its operational discipline, intelligence-focused mission profile, and use of intermittent access patterns across an extended campaign timeline [4]. The actor's toolset consists entirely of publicly available open-source utilities, suggesting a calculated choice that minimizes unique forensic indicators. Two tools were deployed: EarthWorm, a tunneling utility supporting SOCKS5 proxy functionality, port forwarding, and multi-hop network bridging aligned with MITRE ATT&CK techniques T1090 (Proxy) and T1572 (Protocol Tunneling); and ReverseSocks5, which establishes outbound-initiated SOCKS5 tunnels to allow inbound access through the compromised device [4]. The specific EarthWorm binary observed carries SHA-256 hash `e11f69b49b6f2e829454371c31ebf86893f82a042dae3f2faf63dcd84f97a584`,

with artifacts staged to `/var/tmp/linuxap`, `/var/tmp/linuxda`, `/tmp/.c`, and `/tmp/R5` [4]. Network-based indicators include command-and-control infrastructure at `67.206.213[.]86`, `136.0.8[.]48`, `146.70.100[.]69`, and `149.104.66[.]84` [4].

Unit 42 notes tactical overlaps between CL-STA-1132 and previously documented clusters including CL-STA-0046, Volt Typhoon, and APT41, all of which carry China-nexus assessments under prevailing threat intelligence frameworks, though no definitive nation-state attribution has been published [4]. The use of EarthWorm has appeared across multiple campaigns associated with these groups, though tool sharing within state-sponsored ecosystems complicates single-actor attribution. The targeting of Active Directory infrastructure following firewall compromise is consistent with a strategic intelligence-collection mission: a compromised perimeter device provides visibility into all traversing network traffic, while an enumerated Active Directory provides a roadmap to privileged accounts, certificate infrastructure, and high-value data repositories.

Recommendations

Immediate Actions

The most effective immediate mitigation is restricting the User-ID Authentication Portal to trusted internal IP addresses only, eliminating any internet or untrusted network access path to the vulnerable service [1]. Organizations that do not operationally require the Authentication Portal should disable it entirely through the Device > User Identification > Authentication Portal Settings panel. Palo Alto Networks additionally recommends disabling Response Pages within Interface Management Profiles on all untrusted interfaces. For organizations with active Threat Prevention subscriptions running PAN-OS 11.1 or later, enabling Threat ID 510019 provides signature-based detection and blocking for known exploitation attempts against this vulnerability [1].

All organizations with internet-exposed PAN-OS User-ID Authentication Portals should immediately initiate forensic investigation for pre-disclosure compromise. Because the attackers systematically cleared nginx logs, crash records, and core dump files following successful exploitation, an absence of log evidence cannot confirm an absence of intrusion. Security teams should inspect target staging paths (`/var/tmp/linuxap`, `/var/tmp/linuxda`, `/tmp/.c`, `/tmp/R5`), examine network flows for connections to the four published command-and-control IP addresses, and review Active Directory query logs for anomalous activity originating from firewall management interfaces during the April 9 to May 6, 2026 window.

Short-Term Mitigations

Palo Alto Networks' patch releases began on May 13, 2026, with a second wave of hotfixes scheduled for May 28, 2026 [1]. The following table summarizes patched versions by PAN-OS branch:

PAN-OS Branch	Patched Versions — Wave 1 (May 13)	Patched Versions — Wave 2 (May 28)
12.1	12.1.4-h5	12.1.7
11.2	11.2.7-h13, 11.2.10-h6	11.2.4-h17, 11.2.12
11.1	11.1.4-h33, 11.1.6-h32, 11.1.10-h25, 11.1.13-h5	11.1.7-h6, 11.1.15
10.2	10.2.10-h36, 10.2.18-h6	10.2.7-h34, 10.2.13-h21, 10.2.16-h7

Organizations should prioritize patching internet-facing firewalls and devices protecting high-value network segments first, establishing patch deployment timelines that account for operational change-control requirements without allowing unnecessary delay. The official Palo Alto advisory provides the authoritative version-specific patch schedule and should be monitored for any updates.

Strategic Considerations

This exploitation campaign illustrates a structural risk that extends beyond the immediate patch cycle. Network perimeter security devices—firewalls, VPN concentrators, load balancers—are positioned to enforce organizational security boundaries, yet they are complex software systems carrying their own attack surfaces that require continuous assessment. The principle of least exposure applies as forcefully to security infrastructure as to any application server: management interfaces, authentication portals, and diagnostic endpoints on perimeter devices should be accessible only from explicitly authorized sources, with network-level controls that do not depend on the firewall's own policy enforcement to protect the firewall itself.

The campaign's pre-disclosure exploitation window—more than three weeks elapsed between first observed activity and public advisory—underscores the necessity of behavioral detection coverage for network infrastructure devices. Signature-based detection is inherently reactive to published indicators; the organizations that detected this activity during April did so through behavioral anomalies. Security operations programs should ensure that firewall management planes are included in monitoring scope, with alerting on anomalous outbound connections, unexpected process execution, and file system

changes in staging directories. Treating network edge devices as managed, instrumented endpoints rather than trusted black boxes is increasingly necessary as state-sponsored adversaries systematically target them as high-value initial access vectors [4].

CSA Resource Alignment

CVE-2026-0300 and the associated CL-STA-1132 campaign connect directly to several core Cloud Security Alliance frameworks. CSA's Software-Defined Perimeter Architecture Guide V3 articulates an "authenticate-before-connect" architecture that would significantly constrain the attack surface this vulnerability requires [5]. SDP's fundamental design principle—rendering infrastructure invisible and inaccessible to unauthenticated parties before any connection is permitted—is designed to eliminate the exposed portal condition that CVE-2026-0300 requires, effectively constraining this attack surface when correctly deployed. Organizations reassessing their perimeter access control architecture in response to this advisory should consult the SDP guide as a reference for designs that eliminate unauthenticated entry points to network services at an architectural level rather than relying on host-level mitigations.

CSA's Zero Trust Guidance for Achieving Operational Resilience addresses the organizational and architectural practices needed to limit blast radius when perimeter devices are compromised [6]. The guidance emphasizes continuous validation of device posture and network segmentation practices that contain lateral movement from compromised edge devices. The adversary's successful escalation from firewall compromise to Active Directory enumeration and second-device compromise within days is a concrete example of the lateral movement pathways that Zero Trust segmentation disciplines are designed to interrupt. The guidance also addresses identity-centric controls that do not treat network position as a trust signal—a posture directly applicable to environments where the network perimeter device itself is compromised.

The Cloud Adversarial Vectors, Exploits, and Threats (CAVEaT) framework, developed by CSA in collaboration with MITRE, provides a structured taxonomy for analyzing this campaign's kill chain [7]. The adversary's observed progression—from exploitation of a network-facing service, to credential and identity enumeration, to persistent tunneling and multi-device compromise—maps across multiple CAVEaT attack vectors relevant to hybrid cloud and on-premises network environments. Security architects and threat modeling practitioners using CAVEaT should incorporate network edge device compromise as an initial access vector alongside cloud-native attack paths, ensuring that defensive controls are tested against adversary behavior patterns consistent with state-sponsored threat actors rather than only commodity threat profiles.

References

[1] Palo Alto Networks. "CVE-2026-0300 PAN-OS: Unauthenticated user initiated Buffer Overflow Vulnerability in User-ID™ Authentication Portal." Palo Alto Networks Security Advisories, May 2026.

[2] CISA. "Known Exploited Vulnerabilities Catalog." Cybersecurity and Infrastructure Security Agency, May 6, 2026.

[3] BleepingComputer. "Palo Alto Networks warns of firewall RCE zero-day exploited in attacks." BleepingComputer, May 2026.

[4] Palo Alto Networks Unit 42. "Threat Brief: Exploitation of PAN-OS Captive Portal Zero-Day for Unauthenticated Remote Code Execution." Unit 42, May 2026.

[5] Cloud Security Alliance. "Software-Defined Perimeter Architecture Guide V3." Cloud Security Alliance, May 2026.

[6] Cloud Security Alliance. "Zero Trust Guidance for Achieving Operational Resilience." Cloud Security Alliance, April 2026.

[7] Cloud Security Alliance. "Cloud Adversarial Vectors, Exploits, and Threats (CAVEaT™)." Cloud Security Alliance, November 2023.

