


Ivanti Connect Secure Zero-Day Chain: Root on Patched VPNs

Persistence Without a Backdoor, and a Privilege Escalation Flaw Still Unpatched

2026-10-16

 AI-assisted Rapid Research



CSAI

cloud
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

Volexity observed a previously unknown actor, tracked as UTA0576, chain three Ivanti Connect Secure (ICS) vulnerabilities to obtain unauthenticated root code execution on an appliance that was fully patched at the time, running 22.8R2.4, before Ivanti's October 13 fixes existed. Volexity places the start of exploitation no later than October 3, 2026, while CISA-related coverage cites activity from at least October 1 [1][2][4]. Ivanti's October 13 releases address the authentication bypass and remote code execution components (CVE-2026-104935). According to Volexity, the privilege escalation component has no CVE and no fix yet [1][3].

The actor deployed no web shell or backdoor. Commands and results traveled inside the exploit chain's own HTTPS requests, so each action re-exploited the appliance, and the memory-only modifications may not be visible to Ivanti's Integrity Checker Tool [1]. Post-exploitation focused on Active Directory credentials, LAPS passwords, and BitLocker recovery keys, and the payload selectively deleted log entries [1][2].

Patching alone does not close this exposure. Because the integrity tool may not detect memory-only activity, organizations should patch immediately and then treat any ICS appliance that ran an affected version while internet-reachable since at least October 1 as potentially compromised until network and memory forensics indicate otherwise. This is a precautionary judgment and not a finding that every such device was exploited.

Background

Ivanti Connect Secure is a remote-access VPN gateway that sits on the network perimeter and, by design, holds or brokers access to directory credentials. That position makes it a recurring target. Rapid7 notes that ICS has experienced zero-day exploitation in both 2024 and 2025, and the October 2026 chain continues that pattern [2]. BleepingComputer, citing Shadowserver scan data, reports more than 5,500 internet-accessible ICS instances, and a second BleepingComputer article gives a figure of over 5,600. Neither figure is dated, and neither indicates how many of those instances remain vulnerable [3][4].

On October 13, 2026, Ivanti published an advisory for CVE-2026-104935, rated high severity with a CVSS v3.0 score of 8.6 and classified as authentication bypass by spoofing (CWE-290). The flaw lets a remote, unauthenticated attacker execute arbitrary commands [2]. Fixed releases are 22.8R2.5, 25.1.1.2, 25.1.3.0, and 25.1.3.1, while 22.8R2.4 and earlier, 25.1.1.1 and earlier, and 25.1.2.2 and earlier are affected [2][3]. Ivanti acknowledged exploitation in the wild, and CISA added the CVE to its Known Exploited Vulnerabilities (KEV) catalog the same day [2][4]. The Canadian Centre for Cyber Security issued a corresponding advisory, AV26-1035 [6].

Volexity's own report gives the fuller picture. The actor was using the chain against a fully patched appliance, meaning that patch level provided no protection at the time [1]. Reporting on the earliest exploitation date varies slightly: Volexity places the chain's use at no later than October 3, while CISA-related coverage states activity since at least October 1 [1][4]. These are not strictly the same kind of statement, since the first is an upper bound on when exploitation began and the second is a floor on observed activity. Defenders should begin log hunting from October 1, the earlier date, to avoid missing early activity. The technical detail in this note rests largely on Volexity's report, and Ivanti has not published equivalent detail on the third flaw.

Security Analysis

The exploit chain

Volexity describes three stages. The first is an authentication bypass in `welcome.cgi`. Requests to `/dana-na/auth/url_default/welcome.cgi` can carry an `X-DSCGI-PATH` header that routes the request to a protected script such as `wu.cgi`. The appliance authorizes the request according to the URL path but dispatches it to the destination named in the header, which gives an unauthenticated caller access to functions that should require a session [1].

The second stage is remote code execution in `wu.cgi`. The handler accepts an uploaded file, computes an MD5 hash of the filename, and writes the content to `/home/runtime/tmp/<MD5(filename)>`. A race condition exists between creating the file and deleting it after validation fails, and the actor wins that race by sending parallel requests that execute the file before it is removed [1]. Volexity reports that CVE-2026-104935 covers both of these first two flaws [1].

The third stage converts that foothold into root. The `dsinvoked` service trusts configuration data in the DSCache, and the actor patched the `libdsplibs.so` library in memory so that routine maintenance tasks ran attacker-supplied Perl scripts with root privileges [1]. This is the component without a CVE or a fix. Help Net Security's coverage refers to the same service as a privilege escalation in the `dsinvoked` broker [5]. The table below summarizes the three stages and their status.

Stage	Component	Effect	Status as of 2026-10-16
1. Authentication bypass	<code>welcome.cgi</code> via <code>X-DSCGI-PATH</code> header	Unauthenticated access to protected scripts	Fixed in 22.8R2.5, 25.1.1.2, 25.1.3.0, 25.1.3.1 (CVE-2026-104935) [1][2]
2. Remote code execution	<code>wu.cgi</code> upload race condition	Arbitrary command execution	Fixed in the same releases (CVE-2026-104935) [1][2]
3. Privilege escalation	<code>dsinvoked</code> / <code>libdsplibs.so</code> in-memory patch	Command execution as root	No CVE assigned; no fix released [1][3]

Persistence without a backdoor

The subtitle of this note refers to persistence without a backdoor, and the distinction matters for defenders. Volexity reports that UTA0576 "passed commands and results through the exploit chain's HTTPS requests rather than a separate command-and-control (C2) channel, with no backdoor or web shell deployed to the compromised device" [1]. Volexity's account implies that each operation required the chain to be replayed. The actor's durable access depended on the vulnerabilities remaining reachable, not on code left on disk. In that sense the persistence is a property of the unpatched flaws more than of an implant.

This has two consequences. First, applying the October 13 fixes breaks the first two stages and should remove the actor's ability to re-enter through this route, assuming no other access method was established, for example through harvested credentials. It does not undo anything already taken: credentials, LAPS passwords, and BitLocker recovery keys that were exfiltrated remain valid until rotated.

Second, because the privilege escalation stage is still unpatched, an attacker who finds another route to code execution on the appliance could reuse that stage. This is an inference from the reported patch status and not something Volexity or Ivanti has reported, and it should be weighed accordingly.

Anti-forensics and the limits of the Integrity Checker

Volexity found that the payload selectively removed entries from `/home/runtime/dlogs/debuglog` and `/var/log/audit/audit.log` while preserving older records, and that it deleted itself and its temporary files, restored the privilege escalation configuration, and reversed its memory patches through `/proc/self/mem` [1]. Because the modifications existed only in memory, Volexity cautions that the Ivanti Integrity Checker Tool "may not be able to identify device compromise" [1]. Ivanti's guidance to run the tool routinely remains useful, but a clean result should not be read as evidence that an appliance was not exploited [3].

Volexity's memory analysis demonstrated that volatile evidence can survive this cleanup. It recovered a `dsstatmond` process performing packet capture from a hidden `.iscap` directory, along with Python scripts, JSON output, and PCAP files, directly from memory [1]. The practical lesson is that memory acquisition should happen before a suspect appliance is rebooted or upgraded, since either action discards the evidence that the logs no longer contain.

Post-exploitation objectives

The activity Volexity observed was directed at the surrounding Active Directory environment. The actor used custom LDAP and Kerberos tooling to extract credentials, collected LAPS passwords and BitLocker recovery keys, captured traffic to look for plaintext protocols such as LDAP, HTTP, FTP, and telnet, ran password spraying against Active Directory accounts, and reviewed configuration and authentication logs [1]. Rapid7 and Help Net Security describe the same categories of activity [2][5]. A VPN gateway that can authenticate users against a directory is a position from which directory credentials are accessible, and the downstream exposure extends well beyond the appliance.

Attribution and tooling

Volexity assesses with low confidence that the actor has a criminal motivation, based on targeting patterns, and with medium confidence that large language models were involved in developing the exploit and the post-exploitation scripts [1]. These are the vendor's assessments, the evidence behind them is not public in detail, and they have not been independently confirmed. Alongside the compressed patch-to-exploit windows described in CSA's SonicWall notes [7][8], they are one more data point

suggesting that attackers may be shortening the time between disclosure and exploitation of edge devices. CSA's Cisco SD-WAN note describes a related pattern of root access obtained before public disclosure [9].

Points of divergence in reporting

Secondary sources differ in emphasis on a few details. Help Net Security quotes Ivanti as saying it initially discovered a portion of the vulnerability through its product security program and patched it in the July release 25.1.3.0, with backporting under way when exploitation occurred [5]. This is best read as a partial earlier fix and not as a contradiction of the October 13 account, and it helps explain why 25.1.3.0 appears among the fixed releases. The federal remediation deadline is reported as October 16 by Help Net Security and as "Friday" under BOD 26-04 by BleepingComputer, which are consistent with each other [4][5]; defenders should still confirm the date against the CISA KEV entry. As of 2026-10-16 the third flaw has no CVE, so readers should check Ivanti's advisory and Volexity's repository for updated indicators of compromise and any later CVE assignment.

Recommendations

Immediate Actions

Upgrade every ICS appliance to 22.8R2.5, 25.1.1.2, 25.1.3.0, or 25.1.3.1 on an emergency basis outside normal patch cycles, as the KEV listing and Ivanti's advisory direct [2][3]. Before upgrading or rebooting any appliance that was internet-reachable since October 1, capture memory and network traffic so that the in-memory artifacts Volexity describes are preserved [1]. Hunt in web and proxy logs from that date for high volumes of requests to `/dana-na/auth/url_default/welcome.cgi` that carry an `X-DSCGI-PATH` header, particularly where the header path differs from the URL path, and check `/tmp/` on the appliance for post-exploitation artifacts [1]. Prioritize appliances by exposure and by the directory privileges they hold.

If exploitation cannot be ruled out, assume that directory secrets reachable from the appliance may have been taken. That means rotating credentials for service accounts used by the VPN, resetting LAPS passwords, rotating BitLocker recovery keys, and reviewing Active Directory authentication logs for password spraying and anomalous LDAP or Kerberos activity [1]. If compromise is confirmed, follow Ivanti's recovery procedure [2].

Short-Term Mitigations

Because the privilege escalation flaw is unpatched, reduce the appliance's reachable attack surface where operations allow, for example by restricting management interfaces and limiting source networks for the login portal. Forward appliance logs to an external collector in real time, since the payload deletes local entries after the fact [1]. Segment the appliance's network position so that its access to domain controllers is limited to what authentication requires, and monitor that path for bulk LDAP queries. Subscribe to Ivanti's advisory feed and plan an out-of-cycle change window for the privilege escalation fix when it ships [1][3].

Strategic Considerations

This incident shows the limits of treating the vendor's integrity tooling as a compromise detector, and organizations should build independent telemetry for edge devices, including network capture, off-box logging, and a rehearsed memory-acquisition process. Reducing dependence on a single perimeter VPN concentrator may reduce the value of any single device, for example through zero trust access architectures that avoid placing broad directory access on a perimeter appliance [10]. Finally, the repeated exploitation of ICS in 2024, 2025, and 2026 supports including edge-device vendor track record and patch cadence in third-party risk decisions [2]. Vendor history should be one factor among several, since comparable patterns exist across vendors [7][8][9].

CSA Resource Alignment

CSA's most directly related prior work is its set of edge-appliance notes, which describe the same pattern on other vendors' remote-access and network gateways. "SonicWall SMA Zero-Days: Edge Appliance Root Returns" covers chained, actively exploited flaws that gave unauthenticated root compromise, and "SonicWall SMA 1000: Exploitation Attempts Within Two Days of Patch" addresses the compressed window between patch release and exploitation [7][8]. "Cisco SD-WAN CVE-2026-20245 Zero-Day: Root Access Pre-Disclosure Exploitation" documents root access obtained before public disclosure on a network appliance [9]. Read together with this note, they indicate that the response expectations for edge gateways, namely emergency patching, memory-capable forensics, and credential rotation, apply across vendors and not to Ivanti alone.

CSA's executive briefing on how Zero Trust protects organizations by securely connecting users to resources from anywhere speaks to the strategic point above, since it discusses replacing traditional VPN-centered access with architectures that grant narrower, per-resource access [10]. For control

mapping, the AI Controls Matrix v1.1 provides threat and vulnerability management and logging domains that correspond to the patching, off-box log retention, and monitoring recommendations here, although the matrix is oriented to AI systems and applies to this incident only by analogy [11].

References

- [1] Volexity. "[A Series of Unauthenticated Events: Chained Exploitation of Ivanti Connect Secure Zero-Da ys](#)." Volexity Blog, October 13, 2026.
- [2] Rapid7. "[Ivanti Connect Secure zero-day exploited in the wild \(CVE-2026-104935\)](#)." Rapid7 Blog, October 2026.
- [3] BleepingComputer. "[Ivanti warns of Connect Secure zero-day exploited in attacks](#)." BleepingComputer, October 2026.
- [4] BleepingComputer. "[CISA warns admins to immediately fix Ivanti flaw exploited as zero-day](#)." BleepingComputer, October 2026.
- [5] Help Net Security. "[Attackers chained three Ivanti zero-days to gain root on fully patched VPN appliances](#)." Help Net Security, October 14, 2026.
- [6] Canadian Centre for Cyber Security. "[Ivanti security advisory \(AV26-1035\)](#)." Canadian Centre for Cyber Security, October 13, 2026.
- [7] Cloud Security Alliance. "[SonicWall SMA Zero-Days: Edge Appliance Root Returns](#)." CSA Labs, July 2026.
- [8] Cloud Security Alliance. "[SonicWall SMA 1000: Exploitation Attempts Within Two Days of Patch](#)." CSA Labs, October 2026.
- [9] Cloud Security Alliance. "[Cisco SD-WAN CVE-2026-20245 Zero-Day: Root Access Pre-Disclosure Exploitation](#)." CSA Labs, June 2026.
- [10] Cloud Security Alliance. "[An Executive View on How Zero Trust Protects Organizations by Securely Connecting Users to Resources from Anywhere](#)." Cloud Security Alliance, 2025.
- [11] Cloud Security Alliance. "[AI Controls Matrix \(AICM\) v1.1](#)." Cloud Security Alliance, 2026.