Hive Pro

Threat Level

Red

Hiveforce Labs

THREAT ADVISORY

VULNERABILITY REPORT

CVE-2026-6875: ServiceNow AI Platform Pre-Auth RCE Exploited

Date of Publication	Admiralty Code	TA Number
July 21, 2026	A1	TA2026206

Summary

First Seen: April 2026
Affected Products: ServiceNow AI Platform (formerly Now Platform)
Impact: CVE-2026-6875 is a critical flaw in the ServiceNow AI Platform (formerly the Now Platform), a widely used service for building and automating business workflows. The cause is a code injection weakness: an attacker with no login can break out of ServiceNow's script sandbox and run their own code on a targeted instance. The impact is severe, because a successful attack can fully take over the affected instance and any proxy servers connected to it, exposing sensitive data and disrupting critical workflows. The good news is that a fix is available. ServiceNow patched its own hosted instances back in April 2026 and, on July 13, 2026, released updates for self-hosted customers and partners, so upgrading to the latest patched release closes the hole. One point to keep in mind: threat intelligence reports the flaw is now being exploited in the wild, while ServiceNow says it has not confirmed any exploitation of the instances it hosts.

CVE

CVE	NAME	AFFECTED PRODUCT	ZERO-DAY	CISA KEV	PATCH
CVE-2026-6875	ServiceNow AI Platform Remote Code Execution Vulnerability	ServiceNow AI Platform	✗	✗	✓

Vulnerability Details

#1

CVE-2026-6875 is a critical flaw in the ServiceNow AI Platform that hands an unauthenticated attacker the ability to run their own code on a target instance. In plain terms, it is a code injection weakness, tracked as CWE-94 (Improper Control of Generation of Code): a crafted request reaches ServiceNow's server-side script engine and escapes the sandbox that is supposed to contain it, turning what should be harmless input into attacker-controlled code execution on the instance.

#2

The flaw is reachable before authentication. The exploit path targets a pre-authentication endpoint, /assessment_thanks.do, and from there abuses a sandbox-escape gadget to reach a code-execution primitive. This means an attacker who can send requests to an exposed ServiceNow AI Platform instance does not need valid credentials to attempt exploitation.

#3

The vulnerability is described as exploitable only in high-complexity attacks, but with severe consequences when successful: unauthenticated code execution and full compromise of the ServiceNow instance along with any proxy servers connected to it.

#4

The vulnerability affects the ServiceNow AI Platform across multiple release families. ServiceNow addressed it in the following fixed releases: Australia Patch 2, Yokohama Patch 12 Hot Fix 1b, Yokohama Patch 13, Zurich Patch 7b, Zurich Patch 9, Brazil EA, and Brazil GA. Instances running versions before these fixed releases are affected. ServiceNow remediated its own hosted instances starting in April 2026 and made the corresponding security updates available to self-hosted customers and partners on July 13, 2026. The updates additionally ship a new feature called Guarded Script, which restricts the type of code that can run in sandbox contexts and is intended to make future sandbox escapes less likely.

#5

Evidence of active exploitation is contested. A threat intelligence firm reported that the first exploitation attempts appeared on Friday, July 17, 2026, and confirmed active in-the-wild abuse over the following weekend. The observed payloads hit the same pre-authentication sink, but the attackers' sandbox-escape gadget reaches the same code-execution primitive through a different route than the published proof-of-concept. ServiceNow, by contrast, has not flagged the vulnerability as actively abused; its official advisory still states that it is not currently aware of exploitation against ServiceNow instances, and a company spokesperson said that based on its investigation to date it had not observed evidence that the reported activity is related to instances ServiceNow hosts.

Vulnerability

CVE ID	AFFECTED PRODUCTS	AFFECTED CPE	CWE ID
CVE-2026-6875	ServiceNow AI Platform (before Australia Patch 2, Yokohama Patch 12 Hot Fix 1b, Yokohama Patch 13, Zurich Patch 7b, Zurich Patch 9, Brazil EA, and Brazil GA)	cpe:2.3:a:servicenow:ai_platform:*:*:*:*:*:*:*:*	CWE-94

Recommendations

Apply the ServiceNow Security Update Immediately: Self-hosted users and partners should upgrade affected instances to a patched family release without delay. Confirm your instance is running Australia Patch 2, Yokohama Patch 12 Hot Fix 1b, Yokohama Patch 13, Zurich Patch 7b, Zurich Patch 9, Brazil EA, Brazil GA, or later, as these releases contain the fix for CVE-2026-6875.

Prioritize Self-Hosted and Partner Instances: ServiceNow has already remediated the instances it hosts, so the remaining risk is concentrated in self-hosted and partner-managed deployments. Inventory every self-hosted ServiceNow AI Platform instance and treat any that has not received the July 13, 2026 update as a priority for patching.

Hunt for Exploitation Attempts: Review web and application logs for requests targeting the pre-authentication endpoint /assessment_thanks.do, particularly unusual or malformed payloads sent to it. Because in-the-wild attackers are using a route that differs from the published proof-of-concept, focus detection on the targeted endpoint and on post-exploitation signs such as unexpected code execution, new processes, or outbound connections from the instance and any connected proxy servers.

Enable Guarded Script: The security updates introduce Guarded Script, a feature that restricts the type of code allowed to run in sandbox contexts and reduces the likelihood of future sandbox escapes. After patching, enable and validate this hardening feature as an additional layer of defense against this class of flaw.

Reduce External Exposure: Limit direct internet exposure of ServiceNow AI Platform instances where possible, and place them behind appropriate access controls or a web application firewall. Restricting who can reach the vulnerable pre-authentication endpoint lowers the practical opportunity for an unauthenticated attacker to attempt exploitation.

Potential MITRE ATT&CK TTPs

Tactic	Technique	Sub-technique
Initial Access	T1190: Exploit Public-Facing Application	
Execution	T1059: Command and Scripting Interpreter	
Resource Development	T1588: Obtain Capabilities	T1588.006: Vulnerabilities

Patch Link

https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB3137947

References

https://support.servicenow.com/kb?id=kb_article_view&sysparm_article=KB3137947

https://x.com/DefusedCyber/status/2078418391321219448

https://slcyber.io/research-center/smashing-the-servicenow-sandbox-pre-authentication-rce/

What Next?

At **Hive Pro**, it is our mission to detect the most likely threats to your organization and to help you prevent them from happening.

Book a Demo of **HivePro**.

REPORT GENERATED ON

July 21, 2026 • 07:20 AM

More at www.hivepro.com