CSAI Foundation | Cloud Security Alliance

# Three CVSS 10.0 Flaws in ServiceNow's AI Platform

## Unauthenticated RCE and SQL Injection Reach the GenAI Backbone

2026-08-29

AI-assisted Rapid Research

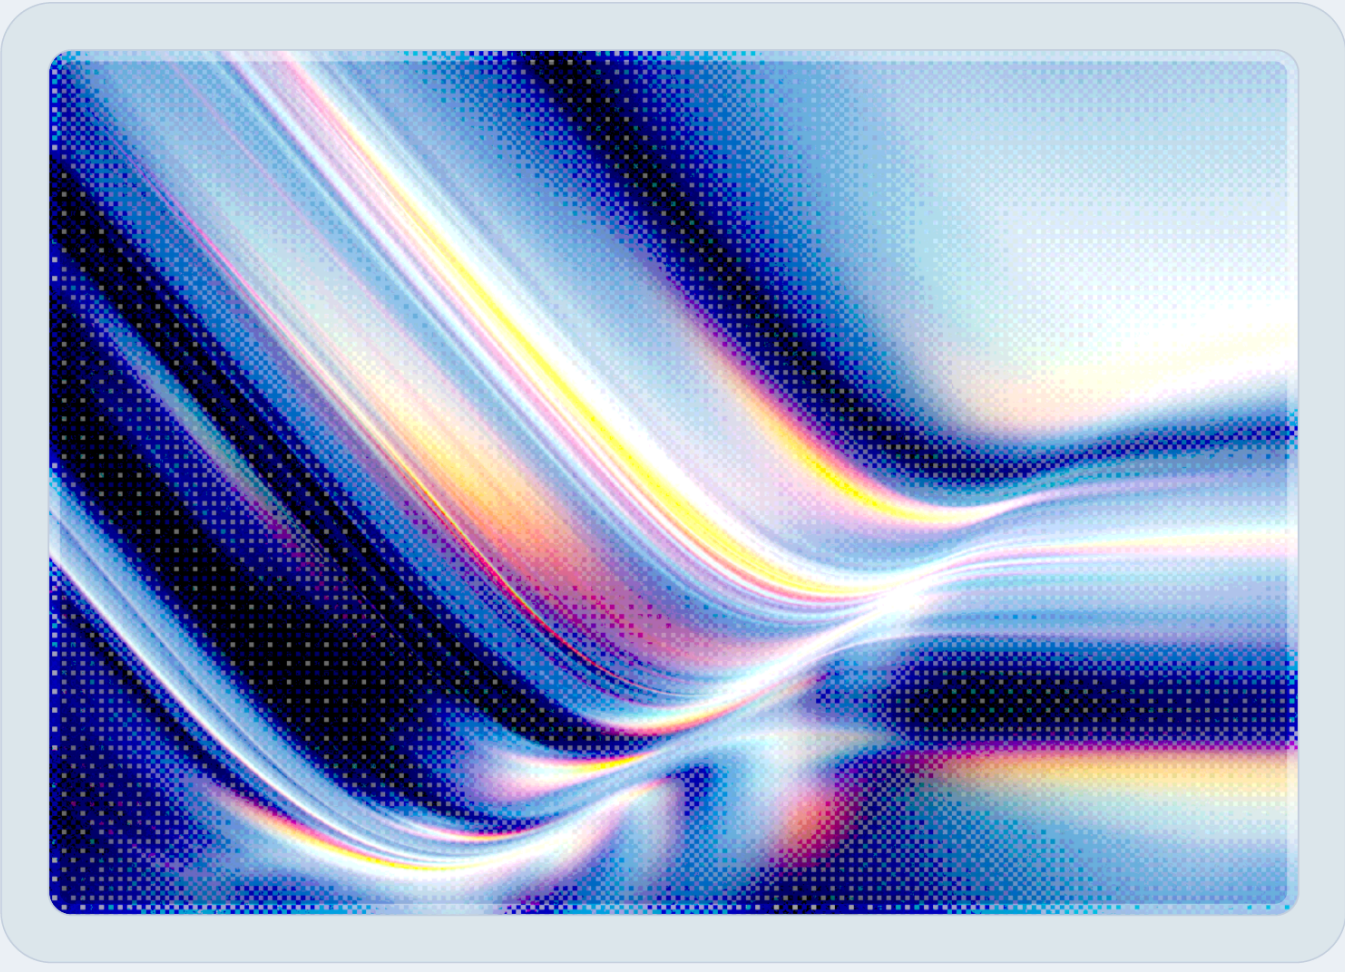

CSAI

CSA cloud security alliance®

## Key Takeaways

- On August 27, 2026, ServiceNow disclosed four vulnerabilities in its Now Platform and ServiceNow AI Platform, three of which received the maximum possible CVSS score of 10.0 and are exploitable without authentication [1][2].
- The most consequential of the three given its role in ServiceNow's AI architecture, CVE-2026-18885, is a code injection flaw in the GraphQL Composite Data API — the same interface ServiceNow markets as the mechanism through which its Now Assist AI Agents read and act on platform data — and it permits an unauthenticated attacker to execute arbitrary code and access or modify instance data [1][3].
- A second maximum-severity flaw, CVE-2026-18886, abuses improper access control in the system configuration image upload processor to let an unauthenticated user create or modify instance data and escalate privileges; a third, CVE-2026-74820, is a SQL injection vulnerability that lets an attacker run arbitrary statements against the instance's underlying database [1][2].
- A fourth flaw, CVE-2026-6876, rated 8.7 and involving a sandbox escape that can also lead to unauthenticated code execution, rounds out a disclosure batch that security commentators described as unusually severe for a single advisory [1][2].
- ServiceNow says it has found no evidence of active exploitation as of disclosure and has already patched its cloud-hosted instances [1][2]. Customers running self-hosted or partner-managed instances must apply the fix themselves; in this analysis, the exposure window for those environments is the primary near-term risk.
- Because the vulnerable component sits underneath ServiceNow's AI Agents and GenAI workflow features, this disclosure adds a third data point to a pattern CSA's own research has now documented twice this year in different AI-adjacent platforms: as vendors wire large language model agents directly into core data and orchestration layers, defects in that plumbing tend to inherit an unauthenticated, pre-authentication, maximum-severity risk profile [4][5]. Three platforms is a pattern worth tracking, not yet proof that every AI infrastructure layer behaves this way.

# Background

ServiceNow operates a widely deployed enterprise SaaS platform, and over the past two years it has repositioned that platform — rebranding portions of it as the "ServiceNow AI Platform" — around agentic AI capabilities marketed under the Now Assist name. A central piece of that repositioning is the GraphQL Composite Data API, which a ServiceNow Community article describes as the interface AI Agents use to determine what data they need and what to do with it across the platform [3]. In other words, GraphQL Composite Data is not a peripheral integration endpoint; it is positioned as the connective tissue between ServiceNow's stored enterprise data and the AI agents that increasingly automate workflows against that data.

On August 27, 2026, ServiceNow published security advisory KB3152242 disclosing four vulnerabilities affecting the Now Platform and ServiceNow AI Platform across the Xanadu, Yokohama, and Zurich release families [2]. Three of the four — CVE-2026-18885, CVE-2026-18886, and CVE-2026-74820 — received CVSS base scores of 10.0, the maximum possible rating, and all three are exploitable by an attacker with no authenticated session at all [1][2]. The fourth, CVE-2026-6876, carries a CVSS score of 8.7 and involves a sandbox escape that can likewise culminate in unauthenticated remote code execution [1]. ServiceNow stated in each advisory record that it is not currently aware of exploitation in the wild, and the company said it has already applied the fix to its cloud-hosted customer instances; the patched builds for self-managed environments are identified as Xanadu Patch 11 Hot Fix 7a, Yokohama Patch 12 Hot Fix 3b and Patch 13 Hot Fix 4, Zurich through Patch 12, and Australia patches 2 through 5 [1].

Security commentators reacted to the disclosure with unusual candor about its severity. David Shipley of Beauceron Security remarked that "you never want to see a 10/10 critical, and you really don't want to see three drop in a row..." [2]. Ensar Seker, CISO at SOCRadar, observed that compromising a platform of ServiceNow's centrality "can therefore be much more significant than compromising a standalone application" [2]. Both comments reflect a structural reality of large SaaS platforms: ServiceNow instances typically aggregate IT service management data, HR records, security incident tickets, and — increasingly — the context and credentials that AI agents draw on to take autonomous action, which means a single pre-authentication RCE path has a blast radius spanning a large share of the workflows the platform touches, up to and including any workflow the compromised credentials or agent identity happen to have access to.

# Security Analysis

The three maximum-severity flaws break down into distinct but complementary attack surfaces, summarized below.

| CVE | CVSS | Vulnerability Type | Affected Component | Impact |
|---|---|---|---|---|
| CVE-2026-18885 | 10.0 | Code injection | GraphQL Composite Data API | Unauthenticated arbitrary code execution; read/modify instance data |
| CVE-2026-18886 | 10.0 | Improper access control | System configuration image upload processor | Unauthenticated creation/modification of instance data; privilege escalation |
| CVE-2026-74820 | 10.0 | SQL injection | Instance database layer | Unauthenticated execution of arbitrary SQL statements |
| CVE-2026-6876 | 8.7 | Sandbox escape | Platform scripting sandbox | Code execution following sandbox breakout |

CVE-2026-18885 is the flaw most directly relevant to ServiceNow's AI ambitions. Because the GraphQL Composite Data API is the channel ServiceNow designed for AI Agents to query and manipulate platform data in a single round trip, a code injection defect in that same API means the attack surface exposed to an unauthenticated network attacker is architecturally identical to the interface the platform's own AI agents rely on for legitimate automation [1][3]. This is not a case of a vulnerability existing "near" the AI layer; it is a vulnerability in the data-access layer the AI layer was purpose-built to use, which is why this analysis treats it as the most consequential of the four disclosures — even though ServiceNow's advisory does not itself rank the three CVSS 10.0 flaws relative to one another.

CVE-2026-18886 takes a different path to a related outcome. By abusing improper access control in the image upload processor used for system configuration, an attacker can create or modify data without authenticating first, and then leverage that foothold to escalate privileges within the instance [2]. Combined with CVE-2026-74820's SQL injection path into the underlying database, an attacker does

not need to rely on any single flaw: code execution, data tampering, privilege escalation, and direct database access are each independently reachable pre-authentication through three separate routes into the same platform. CVE-2026-6876's sandbox escape adds a fourth avenue, though at a somewhat lower severity score than the other three, suggesting it may require a more constrained precondition to reach.

This disclosure extends a pattern CSA's own threat intelligence work has documented twice already in 2026: unauthenticated, maximum- or near-maximum-severity remote code execution vulnerabilities in the infrastructure layers that AI agents depend on, rather than in the AI models themselves. CSA's research on the LiteLLM AI gateway examined a chained pair of CVEs (CVE-2026-42271 and CVE-2026-48710) that researchers assessed as equivalent in practical severity to CVSS 10.0 once combined, and CSA's research on the Langflow AI development platform covered CVE-2026-33017, a CVSS 9.8 unauthenticated RCE that was weaponized within roughly 20 hours of disclosure [4][5]. Both cases documented the same structural risk: platforms that concentrate credentials, data access, and execution capability behind agentic AI features become correspondingly higher-value, higher-severity targets when defects surface in that plumbing. ServiceNow's disclosure extends that pattern from open-source AI infrastructure tooling into a commercial enterprise SaaS platform with broad adoption among large enterprises, which raises the stakes given the breadth and sensitivity of data such platforms typically hold.

ServiceNow's statement that it is not aware of active exploitation should be read as a point-in-time assessment rather than a durable guarantee. Public disclosure of exploitation-ready technical detail — even at the level of CVE description and affected-component naming published here — can accelerate independent reproduction by researchers and attackers alike, and the Langflow case cited above was weaponized in a matter of hours once technical details became public [5]. Organizations should treat the interval between disclosure and their own patch application as an active exposure window, not a grace period.

# Recommendations

## Immediate Actions

Organizations running self-hosted or partner-managed ServiceNow instances should confirm and apply the patched builds identified in KB3152242 — Xanadu Patch 11 Hot Fix 7a, Yokohama Patch 12 Hot Fix 3b or Patch 13 Hot Fix 4, Zurich through Patch 12, or the applicable Australia patch — without waiting for a routine maintenance window, given that all three primary flaws require no authentication to exploit [1]. Security teams should also verify, rather than assume, that cloud-hosted instances have received

ServiceNow's server-side fix, since managed-service patching timelines can lag the advisory date for instances behind delayed upgrade schedules or custom configurations. Instances with internet-facing GraphQL or configuration-upload endpoints warrant the highest priority, and any organization that has enabled Now Assist AI Agents against production data should treat this patch as tied to AI feature availability, not merely platform hygiene.

## Short-Term Mitigations

Once patched, security teams should review instance logs from the weeks preceding disclosure for anomalous GraphQL Composite Data API calls, unexpected configuration image uploads, or SQL error patterns consistent with injection attempts, since ServiceNow's "no known exploitation" statement reflects visibility at disclosure time rather than a retrospective guarantee covering the full pre-patch window. Organizations should also inventory which AI Agents and Now Assist workflows have write access to sensitive modules — HR, security incident, and financial data among them — and confirm that role-based access controls limit what any single compromised agent identity could reach, since the value of a foothold in the GraphQL layer scales with how much an authenticated AI agent identity is itself permitted to touch. Where feasible, restricting network exposure of administrative and configuration-upload endpoints to trusted network ranges reduces the practical reachability of CVE-2026-18886 even where patching is delayed.

## Strategic Considerations

This disclosure reinforces the case for treating the data-access layer beneath enterprise AI agent features as security-critical infrastructure in its own right, subject to the same rigor — threat modeling, penetration testing, and expedited patch SLAs — traditionally reserved for internet-facing production systems, rather than treated as an extension of existing SaaS administrative tooling. Organizations evaluating or expanding agentic AI deployments on any enterprise SaaS platform should build vendor risk assessments that specifically probe how the platform's AI agent layer authenticates to and constrains its own data-access APIs, since this incident demonstrates that a flaw in that shared plumbing can bypass authentication entirely rather than merely exceeding an agent's intended permissions. Finally, procurement and GRC teams should request evidence of a vendor's disclosure-to-patch cadence for AI-adjacent CVEs specifically, given how quickly comparable unauthenticated RCE flaws in AI infrastructure have been weaponized elsewhere this year.

# CSA Resource Alignment

This incident connects directly to two pieces of CSA threat intelligence research published earlier in 2026 that documented the same structural pattern — unauthenticated, maximum-severity remote code execution in the infrastructure layer beneath agentic AI features — in different platforms. CSA's LiteLLM AI Gateway: KEV-Listed Attack Chain Enables Full Takeover analyzed a chained pair of CVEs in an AI gateway that concentrates provider credentials and routing logic, assessed by researchers as equivalent in practical severity to CVSS 10.0 once combined; its core recommendation — that AI infrastructure concentrating data access and execution capability be reclassified and hardened as critical security infrastructure — applies directly to ServiceNow's GraphQL Composite Data API given its designed role as the access layer for Now Assist AI Agents [4]. CSA's CVE-2026-33017: Langflow RCE Exploits Enterprise AI Pipelines documented how an unauthenticated RCE in an AI development platform was weaponized within roughly 20 hours of disclosure and led to credential harvesting, reinforcing the recommendation above that the interval between ServiceNow's disclosure and an organization's own patch deployment be treated as active exposure rather than routine lag [5].

More broadly, the governance dimension of this incident — determining what data and actions an AI agent identity should be permitted to reach even when the underlying API is functioning as designed — falls within the scope of CSA's AI Controls Matrix (AICM) v1.1 [6], particularly its threat and vulnerability management and AI systems security domains, which organizations can use to structure the access-control and vendor-assessment reviews recommended above for any enterprise SaaS platform embedding agentic AI capabilities into its core data layer.

# References

[1] The Hacker News. "Three CVSS 10.0 ServiceNow Flaws Could Let Unauthenticated Attackers Execute Code and SQL." The Hacker News, August 28, 2026.

[2] CSO Online. "ServiceNow patches three maximum severity flaws that could put enterprise data at risk." CSO Online, August 28, 2026.

[3] ServiceNow Community. "Going from Zero to Hero in GraphQL and Agentic AI on the Now Platform." ServiceNow Community, 2026.

[4] Cloud Security Alliance. "LiteLLM AI Gateway: KEV-Listed Attack Chain Enables Full Takeover." CSA, June 17, 2026.

[5] Cloud Security Alliance. "CVE-2026-33017: Langflow RCE Exploits Enterprise AI Pipelines." CSA, July 2, 2026.

[6] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2026.

 