CSAI Foundation | Cloud Security Alliance

Indirect Prompt Injection in Atlassian Rovo Exposes Enterprise Data

Two Disclosed Attack Paths Bypass Jira and Confluence Access Controls

2026-08-11

AI-assisted Rapid Research

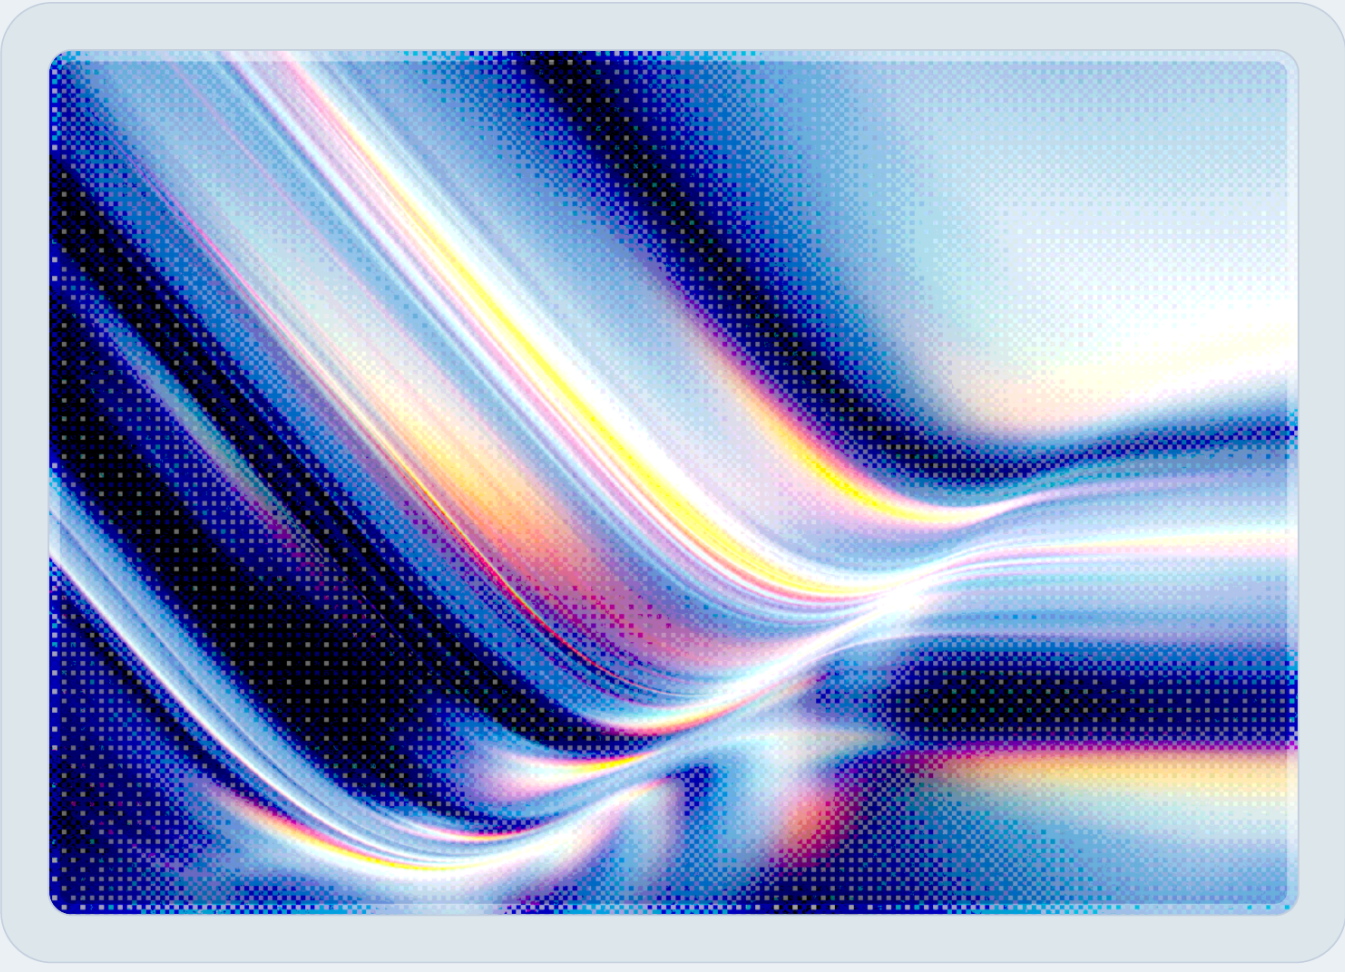

CSAI

CSA cloud security alliance®

Key Takeaways

- Two independent research teams disclosed separate indirect prompt injection paths in Atlassian Rovo, the AI assistant embedded in Jira and Confluence, each capable of exfiltrating data an authenticated user can access to an attacker-controlled destination [1][2][4].
- PromptArmor's content-borne injection hides instructions inside a document a user asks Rovo to process; Rovo's URL retrieval tool then builds a URL containing sensitive data and requests it, delivering the data to the attacker's server logs with no additional user action [2].
- Varonis Threat Labs' "RovoBlast" seeded a malicious prompt directly into a victim's live Rovo session through the unauthenticated `rovoChatPrompt` URL parameter, requiring only a single click and no jailbreak or permission bypass [1][3][4].
- Both attack paths reach data far beyond a single Jira ticket or Confluence page: Rovo's delegated access spans connected systems including SharePoint, Outlook, Slack, Microsoft 365, and Google Workspace, so a successful injection can potentially reach data across whichever of those connectors the compromised session has access to [1][4].
- Disabling Rovo's web search setting does not close the exfiltration path PromptArmor identified, because the underlying tool used to open dynamically constructed URLs remains available even when search is turned off [2].
- Atlassian fixed the RovoBlast URL-parameter flaw server-side by July 8, 2026, but as of PromptArmor's August 5, 2026 publication, the content-borne injection path it reported in May remained unconfirmed as remediated [1][2].

Background

Rovo is Atlassian's AI assistant layer across Jira, Confluence, and Bitbucket, and it is enabled by default for organizations on Standard, Premium, and Enterprise plans [1]. It is designed to help users search across their workspace, summarize tickets and pages, draft content, and answer questions that span multiple Atlassian products and connected third-party applications. To do this usefully, Rovo carries delegated read access to whatever the signed-in user is authorized to see: Jira issues, Confluence pages, and, where an organization has configured them, connectors into SharePoint, Outlook, Slack, Microsoft 365, and Google Workspace [1][4]. Rovo also includes an autonomous ResearchAgent

component capable of multi-source, multi-step web browsing without further user prompting [4]. That combination of broad delegated read access and autonomous outbound browsing reflects a pattern this note and prior CSA research have observed elsewhere in the AI assistant market: it is the structural precondition for prompt injection to become a practical data exfiltration vector rather than a theoretical curiosity [5][6].

Two research teams independently found that this precondition was exploitable in Rovo. PromptArmor disclosed a content-based indirect prompt injection vulnerability to Atlassian on May 23, 2026, and after receiving an initial case-number acknowledgment on May 25 followed by no substantive response through June and July, published its findings publicly on August 5, 2026 [2]. Separately, Varonis Threat Labs disclosed a distinct, URL-parameter-based vulnerability it named RovoBlast through Atlassian's Bugcrowd bug bounty program; Atlassian deployed a server-side fix by July 8, 2026 and paid a $6,000 bounty, and Varonis presented the research publicly at DEF CON 34 in early August 2026 [1][3]. Neither vulnerability has been assigned a CVE identifier as of this writing [1]. The two disclosures describe different technical mechanisms, but both arrive at the same outcome: an AI assistant with legitimate, wide-ranging read access to enterprise data can be redirected by content it was never supposed to trust into sending that data somewhere the organization never authorized.

This pattern is not unique to Atlassian. Enterprise AI assistants integrated into Microsoft 365 Copilot have suffered a comparable sequence of one-click and zero-click exfiltration disclosures over the past year, and CSA's own research has tracked the recurrence of this failure mode as AI assistants are wired into ever more SaaS platforms with real organizational data behind them [5][6]. The Rovo disclosures are best read as the latest instance of a now-familiar architectural problem rather than an isolated flaw in one vendor's product.

Security Analysis

The Content-Borne Injection Path

PromptArmor's research describes a five-step attack chain that begins with entirely ordinary Rovo usage [2]. A user asks Rovo for help organizing or summarizing Jira tickets and, in the course of that request, uploads or references a document that contains concealed prompt injection text. When Rovo processes that document as part of answering the user's request, it treats the embedded text as instructions rather than as inert content to be summarized. Those instructions direct Rovo to gather sensitive information the user's session can access and to construct a URL that encodes that information as parameters, then

to request that URL. Because the destination is an attacker-controlled server, the sensitive data arrives in the attacker's server logs the moment Rovo makes the request; no further action by the attacker or additional click by the victim is required [2].

The vulnerability that makes this possible is narrower than "Rovo can be tricked" in the abstract: PromptArmor identified that Rovo's URL retrieval tool has no protections against opening a URL that has itself been dynamically created by the agent, meaning the tool cannot distinguish between a URL a user intended to visit and one an injected instruction assembled on the fly to carry stolen data [2]. PromptArmor also flagged a second, related weakness: insecure rendering of Markdown image tags, a mechanism the firm describes as a well-known vector for data exfiltration via indirect prompt injection in AI assistants generally. Applied to Rovo, an attacker-controlled image URL with sensitive data appended as parameters would be requested automatically when Rovo renders the response, again without any user action beyond the original, innocuous-seeming request; PromptArmor's writeup frames this as the same class of risk documented in other case studies rather than a chain it independently walked through step-by-step for Rovo [2]. Both weaknesses reflect the same underlying gap — the tool that fetches external content trusts any URL it is given, regardless of whether that URL originated from the user, from Rovo's own reasoning, or from injected instructions buried in a processed document.

Notably, PromptArmor found that this exfiltration path survives one of the more obvious compensating controls an administrator might reach for. Turning off Rovo's web search setting does not remove the underlying tool used to open URLs and retrieve their contents; it only affects how Rovo initiates its own searches. Because the exfiltration relies on that URL-opening capability rather than on search specifically, disabling search leaves the vulnerability fully intact [2]. This Rovo-specific case illustrates a point worth generalizing carefully: administrative toggles that appear to narrow an assistant's capabilities often address only the labeled feature, not the underlying tool call that a prompt injection can invoke directly.

RovoBlast: The URL-Parameter Injection Path

Varonis Threat Labs' RovoBlast attack takes a different route to a similar outcome and requires no document upload at all. Rovo Chat exposes a `rovoChatPrompt` URL parameter intended to pre-fill chat content — a convenience feature that lets a link open Rovo with a suggested prompt already typed in [1][3][4]. Varonis found that this parameter is populated without validating that the content came from a trusted source, so an attacker can craft a link of the form `https://home.atlassian.com/chat?rovoChatPathway=chat&rovoChatPrompt=<attacker-controlled text>` and have it silently seed a victim's live Rovo session the moment the link is clicked, with no warning or indicator that the session has been influenced by an outside party

[1][4]. Varonis also found that the organization-identifying portion of the URL could be omitted entirely, with Atlassian's infrastructure automatically routing the request to the victim's default organization, which removed what might otherwise have been a meaningful barrier to a working exploit link [1].

What turns this into an exfiltration chain, rather than merely an unwanted chat prompt, is Rovo's ResearchAgent, which Varonis found had "almost non-existent" guardrails around instructions received this way [1][4]. ResearchAgent is built to conduct autonomous, multi-source research and to browse and navigate across multiple steps without further user confirmation. Varonis assessed that an injected prompt could plausibly instruct ResearchAgent to retrieve internal content the victim can access across Jira, Confluence, and any connected third-party system, transform that content as needed, and then post it to an external, attacker-reachable destination, within a single automated agent run triggered by one click — a capability the "almost non-existent" guardrails finding makes plausible rather than a chain independently confirmed step-by-step in the public writeup [1][4]. Varonis characterized the vulnerability as requiring no jailbreak and no permission bypass in the traditional sense: the assistant was simply doing what an authenticated user's session was entitled to do, directed by instructions the assistant had no reliable way to identify as untrusted [4].

Why Standard Controls Failed to Catch This

Both disclosures illustrate why conventional application security tooling is poorly matched to this failure mode. Neither attack path requires authentication bypass, privilege escalation, or malware; both operate entirely within the access an already-authenticated, authorized user's session legitimately has. A web application firewall or endpoint detection tool would typically have no reason to flag a signed-in user's browser making a request that Rovo itself initiated on the user's behalf, and a data loss prevention system tuned to catch bulk file downloads would likely not notice a single outbound HTTP request carrying a modest amount of ticket or page content in URL parameters. The vulnerability lives at the boundary between content Rovo is asked to process and instructions Rovo is meant to follow — a boundary that, as in comparable indirect prompt injection cases across the AI assistant market, existing enterprise security stacks were not built to police [5][6].

Recommendations

Immediate Actions

Security teams operating Atlassian Rovo should confirm with Atlassian support whether the content-borne injection path PromptArmor reported has been remediated, since Atlassian's public confirmation lagged the RovoBlast fix by more than a month as of this writing [1][2]. Organizations should also review which user groups and connected applications have Rovo enabled, and disable Rovo access for accounts and connectors — particularly those touching legal, HR, finance, or other highly sensitive Confluence spaces and Jira projects — where the assistant's convenience does not clearly outweigh the exposure [1]. Where feasible, restrict or monitor the connectors Rovo can reach (SharePoint, Outlook, Slack, Microsoft 365, Google Workspace), since each connected system expands the blast radius of a single successful injection [1][4].

Short-Term Mitigations

Because the web-search toggle does not close the content-borne exfiltration path, administrators should not treat it as a compensating control and should instead look for Atlassian guidance on restricting or auditing Rovo's URL-retrieval and image-rendering behavior directly once such guidance is available [2]. Security and IT teams should also instrument and monitor Rovo-initiated outbound requests where visibility permits, since unusual destinations or unusually parameter-heavy URLs in agent-initiated traffic are a meaningful signal that existing SIEM correlation rules for standard web traffic are unlikely to surface on their own [5]. Teams that rely on link-sharing workflows involving Rovo Chat should treat unsolicited or unexpected Rovo links with the same suspicion typically reserved for phishing links, given that RovoBlast required nothing more than a single click on a crafted URL [1][3][4].

Strategic Considerations

The deeper lesson from both disclosures is architectural rather than tactical: an AI assistant that combines broad delegated read access with autonomous outbound browsing or arbitrary URL retrieval creates an exfiltration path that behavioral guardrails and administrative toggles struggle to close completely. Organizations evaluating Rovo or any comparable SaaS AI assistant should treat the separation between trusted control logic and untrusted processed content as a procurement criterion, asking vendors whether URL and image retrieval tools validate the provenance of the URLs they are given, and whether autonomous research or browsing agents can be scoped or gated behind human confirmation before crossing into unfamiliar external destinations. Until such architectural separation is

standard, enterprises should assume that any AI assistant wired into Jira, Confluence, or similar systems with real organizational data behind it is a plausible target for the same content-borne and parameter-borne injection techniques documented here, regardless of vendor.

CSA Resource Alignment

The Rovo disclosures fit a pattern CSA has tracked closely as AI assistants are embedded into SaaS platforms carrying real enterprise data. CSA's research note "Copirate 365: M365 Copilot Command Injection at Scale" examines a structurally similar case — a command injection flaw in Microsoft 365 Copilot in which content embedded in documents, emails, or SharePoint material redirected the assistant's Retrieval-Augmented Generation pipeline into disclosing data it should not have surfaced. That note's observation that Copilot's exposure reflects a systemic pattern across multiple disclosures in a single product line, rather than an isolated bug, applies directly to Rovo: two independently discovered injection paths in the same assistant, disclosed within a matter of months of each other, point to the same conclusion.

CSA's research note "Indirect Prompt Injection Goes Operational" documents the broader 2026 shift of indirect prompt injection from a theoretical concern to an actively exploited technique across agentic AI products, driven by the deployment of agents capable of taking real actions rather than only generating text. Rovo's ResearchAgent — capable of autonomous, multi-step browsing triggered by a single injected prompt — is a strong example of the category of "agent that acts" the note identifies as the precondition for indirect prompt injection to move from theory to operational exploitation. CSA's research note "Agent Data Injection: A New Attack Class Beyond Prompt Injection" extends this analysis to a related but distinct mechanism, in which manipulated or spoofed metadata rather than injected content or URL parameters corrupts an agent's trusted context — underscoring that the untrusted-input problem spans multiple injection surfaces, not only the two documented here [9].

More broadly, both Rovo attack paths are usefully analyzed through CSA's MAESTRO agentic AI threat modeling framework, which addresses the ingestion of untrusted content into an agent's reasoning context across multiple named layers — including input-validation and adversarial-example threats at the model layer and comparable threats at the application and tool-execution layers — spanning the same boundary that both PromptArmor's content-borne injection and Varonis's URL-parameter injection exploit. Organizations formalizing controls for Rovo or comparable assistants should also map those controls to the AI Controls Matrix (AICM) v1.1, whose application and interface security, and threat and vulnerability management domains cover the categories of untrusted-input validation and tool-permission scoping most directly relevant to both disclosed attack paths.

References

[1] Sergiu Gatlan / The Hacker News. "Atlassian Rovo Can Be Tricked Into Sending Jira and Confluence Data to Attackers." The Hacker News, August 2026.

[2] PromptArmor. "Atlassian Rovo Exfiltrates Data, Bypassing Controls." PromptArmor Research, August 5, 2026.

[3] Ryan Naraine / SecurityWeek. "Critical One-Click Vulnerability in Atlassian's Rovo AI Exposed Enterprise Data." SecurityWeek, August 2026.

[4] Varonis Threat Labs. "RovoBlast: How One Click Triggered Atlassian's AI Assistant to Leak Data." Varonis, August 2026.

[5] Cloud Security Alliance AI Safety Initiative. "Copirate 365: M365 Copilot Command Injection at Scale." CSA Research, May 2026.

[6] Cloud Security Alliance AI Safety Initiative. "Indirect Prompt Injection Goes Operational." CSA Research, 2026.

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

[8] Cloud Security Alliance. "MAESTRO: Agentic AI Threat Modeling Framework." Cloud Security Alliance, February 2025.

[9] Cloud Security Alliance. "Agent Data Injection: A New Attack Class Beyond Prompt Injection." Cloud Security Alliance, July 2026.

