

Spam Campaign Abuses Atlassian Jira, Targets Government and Corporate Entities

 trendmicro.com/en_gb/research/26/b/spam-campaign-abuses-atlassian-jira.html

February 17, 2026



Create a site

Sites are the shared space where people organize teams, work, and projects.

Your site

notathreatactor|

.atlassian.net



This site name is just a suggestion. Feel free to change to something your team will recognize.

Continue

Spam

We uncover how a campaign used Atlassian Jira Cloud to launch automated and targeted spam campaigns, exploiting trusted SaaS workflows to bypass security controls.

By: TrendAI™ Research Read time: 7 min (1814 words)

Key takeaways

- Attackers abused Atlassian Cloud's trusted domain for a spate of spam campaigns. The campaigns tried to leverage the domain name and reputation of this legitimate and well-known SaaS platform.
- Emails were tailored to target specific language groups, targeting English, French, German, Italian, Portuguese, and Russian speakers — including highly skilled Russian professionals living abroad.
- These campaigns not only distributed generic spam, but also specifically targeted sectors such as government and corporate entities.
- Keitaro Traffic Distribution System (TDS) powered redirects, channelling targets to dubious investment schemes and online casinos, indicating that financial gain was likely the primary motive behind these campaigns.
- Organisations using Atlassian Jira were prime targets, especially those with high email volume and have a heavy reliance on collaboration tools, environments where Jira notifications are routinely trusted.
- Enterprises should deploy advanced email security solutions such as TrendAI Vision One™ Email and Collaboration Security, which provide layered detection and identity-aware controls to better detect and block phishing and abuse of trusted SaaS platforms.

Introduction

Threat actors used Atlassian Jira Cloud and its connected email system to run automated spam campaigns, effectively bypassing traditional email security by abusing the strong domain reputation of Atlassian Jira Cloud products. The campaigns were active from late December 2025 through late January 2026, during which organisations and individuals worldwide — particularly English, French, German, Italian, Portuguese, and Russian-speaking targets — received spam emails from legitimate-looking Atlassian Jira Cloud addresses.

In addition, campaigns did not appear to generate generic spam. They also targeted specific sectors, most notably government and corporate entities. The emails redirected targets to pages on investment scams and online casino landing sites, suggesting that actors were likely motivated by financial gain.

This activity exemplifies how threat actors can abuse legitimate tools for malicious activities, this time using software-as-a-service (SaaS) platforms to deliver spam or spear-phishing emails. By operating through established and reputable cloud services with strong domain reputations, attackers are able to bypass blocklists and exploit inherent trust in enterprise tools. Traditional email security places higher trust on notifications from SaaS providers. These campaigns piggyback on this trust to exploit a well-known and legitimate SaaS provider. In addition, the built-in compliance with authentication checks like Sender Policy Framework (SPF) and DomainKeys Identified Mail (DKIM) further reduced the likelihood of detection.

The campaign also demonstrates a high degree of automation: Threat actors appear to rapidly create multiple Atlassian instances, likely using free or trial Atlassian accounts, to target specific industry sectors.

As SaaS platforms continue to expand their email-driven workflows, these attacks reveal the need to reassess long-standing trust assumptions and tighten controls around third-party cloud-generated email. In this blog entry, we break down the attack in detail to identify where defences would be the most effective against similar tactics.

We shared this information in advance with Atlassian's security team to report the platform abuse.

Detailed analysis

In this section, we examine how the campaign was executed end-to-end, reproducing key steps to validate our assumptions about the attackers' methods.

Stage 1: Infrastructure provisioning

Threat actors began by creating Atlassian Cloud accounts using randomised naming conventions, enabling them to generate disposable Jira Cloud instances at scale. Our analysis showed that these instances resolved to AWS IP (13[.]227[.]180[.]4), which also served many legitimate Atlassian Cloud deployments. This strongly suggests that the attack used legitimate Atlassian Cloud infrastructure, not compromised servers.

The spam-related instances appeared to have been provisioned without any domain ownership verification. We found no registered domains that corresponded with the instance names. Therefore, the registered domains did not resemble the spoofed company names. This indicates that the threat actors were not attempting to reinforce legitimacy through domain registration but instead relied on the inherent trust associated with Atlassian-generated system emails.

Further review revealed how setting up an Atlassian trial account is a straightforward process, meaning it lowered the barrier for threat actors to repeatedly create new instances.

Figure 1. Creating a trial Jira instance

[download](#)

Figure 2. Jira Kanban Board allows for the creation of automation rules

[download](#)

We also found that Jira offers the ability to bulk-add users via a CSV file. However, doing this would send notification emails to the targets, which could raise suspicion. Therefore, in this case, the attackers used Jira Automation to deliver crafted emails through an integrated email sending platform, which allowed them better control and avoiding alerts to their targets.

We also reproduced the delivery flow of one of the emails observed in the campaigns using Jira Automation rules. In this setup, the only indication of a Jira connection is the sender email address, with nothing in the email body and structure revealing its Jira origin. Notably, the recipient does not need to be a listed user within the instance, nor do they need to accept an invitation to join any project. This allowed the attackers to deliver emails widely and anonymously, without exposing their infrastructure or requiring user enrolment.

Stage 2: Target selection and email generation

TrendAI™ telemetry indicates that attackers targeted recipients with clear intent, focusing on specific organisations, languages, and demographic profiles. In some cases, target lists included highly skilled individuals born in Russia but who are currently living and working abroad, suggesting the campaign had targeted goals, even though financial gain still appeared to be the most prominent objective.

Other target lists were predominantly composed of English, French, German, Italian, and Portuguese-speaking individuals, who received casino-centred spam. The attackers tailored their subject lines to match each recipient's native language. This extra step of crafting tailored subjects demonstrates a deliberate effort to personalise spam content and improve engagement across different regional targets.

Figure 3. Breakdown of targets by industry
[download](#)

Stage 3: Email delivery

For the delivery phase, the attackers relied on an integrated email sending platform to transmit their messages. These emails were able to bypass many initial security controls due to two key factors:

- The use of a **legitimate atlassian.net sender domain**, which inherently benefits from strong domain reputation
- **Valid SPF and DKIM authentication** applied through Atlassian's integrated email system, allowing the messages to appear trustworthy to email filters

The subject lines used in these campaigns were highly localised and tailored to recipient language profiles. Examples include:

Russian

- Application No.<random number> Your confirmation is required.
- *English: Application No.<random number> Your confirmation is required.*
- Application No.<random number> Your confirmation is necessary.
- *English: Application No.<random number> Your confirmation is necessary.*

Italian

- A new gaming opportunity is available
- *English: A new gaming opportunity is available*
- The game is now available
- *English: The game is now available*
- Platform update and offer: №<random number>
- *English: Platform update and offer: №<random number>*

English

- A new year gift from us: id<random number>
- A holiday note for you: id<random number>
- Notification №<random number>. Haven't you picked up your gifts yet? Hurry up.
- Get your first bonus now: id<random number>
- Special Gaming Opportunity: No<random number>

In some cases, the threat actors used standard Jira-generated subject lines, which are less effective in enticing recipients to click on the links associated with online casinos and dubious investment schemes. It is unclear why threat actors used standard Jira subject lines; it might just have been the result of human error or misconfigured automation rules.

Figure 4 shows an email that was sent to one of the targets that was quarantined.

Figure 4. Quarantined spam email that used a Jira-generated subject line

[download](#)

We were able to source similar emails on VirusTotal, which revealed that the lure was early access to an investment service. The email body was written in Cyrillic and the currency referenced was the ruble, further indicating that the campaign specifically targeted Russian-speaking recipients.

Figure 5. Example of an email in Cyrillic, informing the recipient they were selected for an exclusive investment service

[download](#)

Stage 4: User interaction and payload delivery

Based on embedded links found in our telemetry, the spam links directed recipients to several domains, including:

- adrinal[.]com
- barankinyserialxud[.]online
- archicad3d[.]com

We provide a complete list in the indicators of compromise (IOCs) section.

Many emails also contained links pointing to hostname go[.]sparkpostmail1[.]com, which served as an intermediary redirect to the final landing pages. This hostname is tied to an email delivery platform that is often used for legitimate marketing and transactional email. However, as demonstrated in this campaign, it can also be abused by spammers to mask malicious redirects within unsolicited emails.

The URLs in the emails also pointed to a Traffic Distribution System (TDS) called Keitaro. Keitaro is another legitimate product meant for tracking affiliate and advertising campaigns. In this instance, it was weaponised to redirect spam URLs to final landing pages, the content of which ranged from dubious investment schemes and online casinos, with a blurred example shown in Figure 6.

Figure 6. One of the final landing pages of the spam campaign
[download](#)

Targeting analysis

Our analysis indicates that organisations already using Atlassian Jira were among the primary targets. All of the domains we have identified in the campaign had active Atlassian instances, suggesting that threat actors deliberately selected organisations familiar with Jira-related emails to increase the likelihood of a successful delivery. Sectors characterised by high email volume and heavy adoption of collaboration tools might also have been considered as good targets for this campaign, as they would likely trust and routinely interact with Jira-generated notifications.

Proactive security with TrendAI Vision One™

[TrendAI Vision One™](#) is the industry-leading AI cybersecurity platform that centralises cyber risk exposure management, security operations, and robust layered protection.

TrendAI Vision One™ Email and Collaboration Security

To help defend against new spam tactics and ever-evolving cyberthreats, organisations need more than siloed tools. They need a **unified, AI-powered, enterprise cybersecurity platform** that secures users, data, and communication across all layers.

[Email and Collaboration Security](#) delivers exactly that: safe communication and seamless collaboration through AI-powered threat detection and human risk management. The cloud-native solution proactively mitigates risks from AI-generated phishing, business email compromise (BEC), ransomware, and other sophisticated attacks using machine learning, behavioural analysis, and risk-based policy enforcement. This allows organisations to detect threats in real time and maintain their resilience against similar spam campaigns.

TrendAI Vision One™ Threat Intelligence Hub

[TrendAI Vision One™ Threat Intelligence Hub](#) provides the latest insights on emerging threats and threat actors, exclusive strategic reports from TrendAI™ Research, and TrendAI Vision One™ Threat Intelligence Feed in the TrendAI Vision One™ platform.

Emerging Threats: [Atlassian Jira Infrastructure Abuse to Bypass Security Controls](#)

TrendAI Vision One™ Intelligence Reports (IOC Sweeping)

[Atlassian Jira Infrastructure Abuse to Bypass Security Controls](#)

Hunting queries

TrendAI Vision One™ Search App

TrendAI Vision One™ customers can use the Search App to match or hunt the malicious indicators mentioned in this blog post with data in their environment.

| mailMsgSubject:"%Request No." AND mailMsgSubject:"Your confirmation is required"

More hunting queries are available for TrendAI Vision One™ with Threat Intelligence Hub entitlement enabled.

Indicators of compromise (IOCs)

The indicators of compromise for this entry can be found [here](#).

Tags

[Spam](#) | [Articles, News, Reports](#) | [Research](#)

Copyright ©2026 Trend Micro Incorporated. All rights reserved.

This website uses cookies for website functionality, traffic analytics, personalization, social media functionality and advertising. Our Cookie Notice provides more information and explains how to amend your cookie settings. [Learn more](#)