


AdFection: Mercenary Spyware Through Programmatic Ads

Security Implications and Guidance

2026-10-19

 AI-assisted Rapid Research



CSAI

cloud
CSA security
alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

On October 15, 2026, Citizen Lab published what it describes as the first public, technical confirmation that mercenary spyware, meaning hacking tools sold to governments, has reached real devices through targeted online advertisements [1][2]. The technique, which Citizen Lab calls AdFection, delivered spyware associated with the commercial surveillance vendor Candiru to at least 45 devices at the Al Jazeera network, and it required no click or other action from the victim [1][3].

The delivery path ran through ordinary programmatic advertising. An Israeli firm named Discover-Tech bought ad placements on major exchanges, and the ads carried code that reported device identifiers to Discover-Tech infrastructure, which decided whether a given device was a target [2][3]. This design moves the targeting decision out of the ad and into a server, so a reviewer who inspects the ad creative would likely see little that is malicious.

Security teams should treat the ad supply chain as a delivery channel for targeted attacks, particularly for high-risk users such as journalists, government staff, and executives. Practical controls include DNS and network-level blocking of ad and tracking traffic, reduced advertising-identifier exposure on managed devices, timely patching, and retention of network logs for retrospective hunting against the published indicators. Public reporting rounds some figures differently across outlets, and this note flags those differences where they matter.

Background

Candiru is an Israeli commercial surveillance vendor that the US Commerce Department added to its Entity List in November 2021 [4]. Press reports describe the company as now operating under the name Integrity Labs after an acquisition by the US-based Integrity Partners; the reports differ on the exact corporate naming, so readers should treat the ownership details as reported rather than independently confirmed [2][4]. Al Jazeera journalists have been a recurring surveillance target, and one outlet recalls that the network was previously hit with NSO Group's Pegasus spyware in 2020 through a different method, a detail we could not confirm in the sources we reviewed [2].

Citizen Lab's report, titled "How Discover-Tech Turned Targeted Ads into a Pipeline for Spyware," attributes the Al Jazeera infections, which occurred during 2024 and 2025, with high confidence to malicious ads purchased by Discover-Tech [1][4]. Confiant, an ad-security firm, also described a separate technical investigation of the activity [4]. Citizen Lab describes Discover-Tech as an intermediary that served multiple surveillance vendors, and one outlet characterizes it as a bridge between mercenary spyware companies and the advertising ecosystem [1][3]. The researchers examined nearly four billion Discover-Tech bidding records from March through June 2026 [1][4].

Programmatic advertising is the context that makes this possible. In real-time bidding, an auction runs in milliseconds each time a page or app loads an ad slot, and the winning bidder's creative, which often includes third-party JavaScript, executes on the user's device. Bid requests commonly carry an advertising identifier and an IP address, which an advertiser can use to recognize a specific device. An attacker who can win auctions on a small number of targeted devices therefore has a delivery mechanism that does not depend on phishing or on compromising a site the victim visits. Reporting also notes that a similar technique, referred to as "Aladdin," was tied to Predator spyware in late 2025, which suggests this approach is not unique to one vendor [3]. We have not independently verified that earlier case.

Security Analysis

Attack Chain

According to the published reporting, the chain has five steps [2][3]. Discover-Tech buys ad placements through ad platforms. The ad contains callback code, which on rendering sends the device's advertising identifier and IP address to a Discover-Tech server. That server compares the values with a target list; devices that do not match receive nothing further. For matching devices, the server reportedly returns exploit code that runs in the browser or app context. Successfully infected devices then connect to Candiru validation infrastructure and subsequently to command-and-control servers. Citizen Lab describes the observed connection sequence as a Discover-Tech ad server contacting a Candiru validation server and then a command-and-control server, and this sequence is the basis for the hunting guidance discussed below [1][2].

This server-side filtering likely explains why the channel is difficult to detect. Devices outside the target list would receive a benign response, so ad-scanning vendors that sample creatives from those devices would see little that is malicious. Only the intended victims receive a payload, which limits the number of copies of the exploit available to defenders. The reporting indicates that Citizen Lab did not recover the

exploit code, and none of the sources we reviewed name a CVE, so the specific browser or operating-system vulnerabilities used remain unknown [2][3]. Defenders should therefore not assume that a particular patch level protects against this chain.

Scale and Exposure

The reported scale of the ad delivery infrastructure is substantial, but it should be read as exposure rather than infection. The bidding records show Discover-Tech activity, not compromised devices, and the confirmed infections are limited to at least 45 Al Jazeera devices [1][3]. The table below summarizes the figures as outlets report them.

Metric	Reported figure	Source
Confirmed infected devices	At least 45 (Al Jazeera)	[1][2] [3]
Bidding records, March–June 2026	"Nearly four billion" (3.77 billion bids in the more precise account)	[1][3] [4]
Auctions won	442 million	[1][3]
Apps and websites where ads appeared	75,373 apps and 120,768 websites; "more than 120,000 websites and 70,000 mobile apps" in another account	[1][3] [4]

These figures appear to reflect the same underlying counts at different levels of rounding or scope; the app and website totals sum to roughly 196,000, a figure that is our own arithmetic rather than one attributed to a source. Readers who need exact figures should consult Citizen Lab's publication directly [1]. Citizen Lab also reports impressions on networks associated with NATO, UN institutions, European Union bodies, and government agencies in Denmark, Norway, and Sweden, and the ads reached users across dozens of countries [1][4]. This indicates that targeted individuals within sensitive organizations could be reached even where the organization itself was not the target, although the published reporting does not show infections outside Al Jazeera.

Why the Ad Ecosystem Is Exposed

In our reading, three structural properties of programmatic advertising may have contributed to the campaign. First, a bidder can gain reach into tens of thousands of apps and sites, and the reporting does not describe any vetting that prevented this. Second, creatives frequently execute third-party code, so

the ad slot works as a script-execution surface inside the browser or app webview. Third, bid requests carry stable device identifiers, which gives an attacker a means of selecting victims. These observations are inferences from the campaign's design rather than findings stated by Citizen Lab.

Platform Response

Reporting indicates that after Citizen Lab shared its findings, several parties removed Discover-Tech or restricted its access in early October 2026. The two outlets that list these parties name the same five, and they differ on the specific action each took, which is variously described as removal, suspension, pausing for review, or account blocking [2][3]. IAB Europe is a trade association rather than an ad platform, so readers should not assume its action resembled a platform's removal of a buyer; we have not established what it did. The table below records who is named and who did not respond.

Party	Reported status	Source
Google	Named as having removed or restricted access	[2][3]
Beeswax	Named as having removed or restricted access	[2][3]
Vungle	Named as having removed or restricted access	[2][3]
Eskimi	Named as having removed or restricted access	[2][3]
IAB Europe	Named as having acted (trade body; nature of action not established)	[2][3]
Yandex	Did not respond to an inquiry	[2]

The removals reduce exposure to this specific buyer, but they do not change the incentive or the mechanism, and a different intermediary could repeat the approach.

Recommendations

Immediate Actions

Organizations should use the published indicators to search historical logs. The reported indicators include the domain `mtrtb[.]com` associated with Discover-Tech and the SHA-256 hash `255869de85e2a171993fc5eb8a556d873a1b8966e040f6f55926f2fa2d595cc8`, which Citizen Lab

labels as Candiru malware; both match across Citizen Lab's report and secondary coverage [1][2]. Because the connection pattern (ad server, then validation server, then command-and-control server) is more informative than any single domain, hunt teams should look for that sequence in DNS, proxy, and NetFlow data rather than relying on a single hit. Organizations with at-risk users should preserve network and endpoint evidence if any match is found, and should escalate to specialists who can perform mobile forensics [4].

Short-Term Mitigations

Blocking advertising and tracking traffic at the network gateway is a direct control because it requires no changes on the endpoint, and eSecurityPlanet's coverage recommends it together with trusted ad blockers [4]. For managed mobile fleets, mobile device management can limit personalized ad tracking, which reduces the value of the advertising identifier as a targeting key [4]. Devices should stay on current operating system and browser versions, although, as noted above, the exploited vulnerabilities are not public, so patching reduces but does not remove the risk [2][4]. For high-risk users such as journalists, executives, and officials, organizations should consider platform high-security modes such as Apple's Lockdown Mode or Android's Advanced Protection. This is our recommendation and is not drawn from the cited reports, and we have not verified that these modes block this specific chain.

Strategic Considerations

Security leaders should add the ad supply chain to threat models for high-risk populations, alongside email, messaging, and web browsing. Procurement and policy teams that buy advertising should ask their demand-side platforms how new buyers are vetted and whether creatives with server-side conditional behavior can be detected. Organizations that operate or rely on ad exchanges may wish to evaluate stronger buyer verification and creative-execution restrictions regardless of whether regulatory action follows. The 2021 Entity List designation of Candiru [4] did not prevent this campaign. We infer that vendor-focused controls alone are insufficient, since intermediaries sit between vendors and targets, although the infections occurred in 2024 and 2025 and the Entity List's effect on the intermediary itself is not directly tested by this case.

CSA Resource Alignment

CSA's [Security Guidance for Critical Areas of Mobile Computing](#) addresses device management and mobile application distribution [5]. It is the natural reference for the mobile device management recommendation above, since limiting ad tracking and enforcing patch levels on managed fleets are

device-management decisions.

CSA's [Mobile Top Threats](#) report analyzes the principal threats to mobile computing based on expert survey data [6]. AdFection is relevant because it reaches the device through the ad and web content layer rather than through a malicious app, and organizations using that report to prioritize mobile risk should consider whether ad-delivered exploitation is represented in their threat models.

CSA's [Mobile Application Security Testing – Sum-Up & Landscape Overview](#) surveys mobile application security testing tools and practices [7]. Its relevance here is that ad SDKs and webviews embedded in apps are part of an app's attack surface, and testing programs that cover third-party SDK behavior can help assess how much untrusted ad content an application will render.

The [AI Controls Matrix \(AICM\) v1.1](#) provides a general framework for control assurance, and organizations can use its threat and vulnerability management and supply chain domains as a reference point when assessing third-party content dependencies [8]. We note that AICM is oriented toward AI systems, so its application to ad-delivered spyware is by analogy rather than direct coverage. CSA has not, to our knowledge, published guidance specifically on programmatic advertising as an attack surface, and that is a gap worth considering for future work.

References

- [1] Citizen Lab. "[How Discover-Tech Turned Targeted Ads into a Pipeline for Spyware.](#)" Citizen Lab, University of Toronto, October 15, 2026.
- [2] The Hacker News. "[Citizen Lab Traces Zero-Click Candiru Spyware Attacks to Targeted Ads.](#)" The Hacker News, October 2026.
- [3] BleepingComputer. "[Zero-Click Attacks Used Targeted Ads to Deliver Candiru Spyware.](#)" BleepingComputer, October 2026.
- [4] eSecurityPlanet. "[AdFection Zero-Click Spyware Attack Hits Al Jazeera.](#)" eSecurityPlanet, October 2026.
- [5] Cloud Security Alliance. "[Security Guidance for Critical Areas of Mobile Computing.](#)" Cloud Security Alliance.
- [6] Cloud Security Alliance. "[Mobile Top Threats.](#)" Cloud Security Alliance, 2024.
- [7] Cloud Security Alliance. "[Mobile Application Security Testing – Sum-Up & Landscape Overview.](#)" Cloud Security Alliance, 2024.
- [8] Cloud Security Alliance. "[AI Controls Matrix v1.1.](#)" Cloud Security Alliance.