


The Ad Delivery Trust Boundary

When the message is approved but the route to the audience is not

What is a media placement worth if neither the advertiser nor the publisher can fully account for the code and intermediaries that reach the audience?



In this briefing

Special Intelligence Brief · 15 October 2026 · Media Strategy / Reputation / Cyber Risk

—	The Strategic Question & Executive Perspective
01	The October 15 Trigger — Verified Evidence
02	Evidence Discipline — Three Claim Levels
03	Figure 01 — The Two-Track Media Trust Boundary
04	Artefact91 Analysis — Four Strategic Fault Lines
05	The Artefact91 Media Trust Boundary Test
06	Responsibility Without Blame by Association
07	The First 72 Hours — A Decision Protocol
08	Decision Intelligence & Board Questions
—	What to Watch Next
—	Sources & Research Note

HOW TO READ IT

A special briefing, published outside the normal twice-monthly cadence: the executive perspective carries the argument, the numbered sections carry the evidence and the method, and every external figure is attributed to its source.

THE STRATEGIC QUESTION

What is the placement worth?

THE STRATEGIC QUESTION

WHAT IS A MEDIA PLACEMENT WORTH IF NEITHER THE ADVERTISER NOR THE PUBLISHER CAN FULLY ACCOUNT FOR THE CODE AND INTERMEDIARIES THAT REACH THE AUDIENCE?

Executive Perspective

Advertising is commonly governed as a commercial communication: who speaks, what claim is made, which audience sees it and where it runs. New research published today establishes a second question. A programmatic advertisement can carry code and targeting signals used to identify—and in specific documented cases compromise—a person without any click. The message can be innocuous while the delivery route is not. ^{[1][2]}

This is not evidence that conventional programmatic advertising is generally malicious, nor that a brand buying media caused these incidents. It does mean media assurance cannot stop at brand safety, placement verification, reach, fraud or routine privacy compliance. For institutions, campaigns and communications agencies, delivery infrastructure is now a stakeholder-trust question.

ARTEFACT91 POSITION

MEDIA INTEGRITY IS A JOINT COMMUNICATIONS–SECURITY DECISION, WITH A VERIFIABLE CHAIN FROM BUYING INSTRUCTIONS TO EXECUTED CODE.

The critical question is not simply whether the message appeared in the right place, but whether its route to the audience met an acceptable trust boundary.

SECTION 01

The October 15 Trigger — Verified Evidence

Citizen Lab at the University of Toronto published *Exposing AdFaction* on October 15, describing the first public technical confirmation of zero-click mercenary spyware delivered through targeted online ads. Researchers attribute the technique to advertising infrastructure run by Discover-Tech and spyware associated with Candiru, which has also used the name Integrity Labs. ^[1]

AL JAZEERA — Citizen Lab connects targeted ad delivery to Candiru activity through network logs and an extracted malware sample. It reports that at least 45 devices were targeted during 2024–2025 and identifies infection evidence. The figure is a **TARGETED**-device count; it must not be presented as the confirmed number of infected devices. ^[1]

SCALE — Citizen Lab analyzed approximately four billion ad-bidding records from a four-month period and found distribution across nearly 200,000 combined websites and mobile apps. Bid requests, ads delivered, profiles gathered, target devices and confirmed infections have different denominators. Confiant separately estimates roughly ten million ad impressions per day through Discover-Tech; this is an analytical estimate, not a census of hacked devices. ^{[1][2]}

SENSITIVE NETWORKS — Researchers also found ad activity on networks associated with sensitive institutions. An ad appearing on an organizational network does not by itself prove that an individual was selected for surveillance or that a device was compromised. Confiant warns that network identifiers cannot reliably identify individual people. ^[2]

PLATFORM RESPONSE — Following the investigation, Google terminated the company's exchange access. Citizen Lab also reports suspensions or blocks by IAB Europe, Beeswax, Vungle/Liftoff and Eskimi. Discover-Tech had previously appeared in approved or certified vendor registries. Approval status therefore cannot be treated as sufficient evidence of continuing security assurance. ^{[1][3]}

SECTION 02


Evidence Discipline — Three Claim Levels

- 1 **ESTABLISHED** — Researchers linked targeted advertising activity to Candiru spyware at Al Jazeera and documented the technical delivery pathway. Multiple vendor access removals have also been reported. [1][3]
- 2 **OBSERVED, NOT PROOF OF INFECTION** — Wide ad distribution, appearances on institutional networks, device identifiers and profiling capabilities show potential exposure paths. They do not show that every ad contained an exploit, every device was deliberately targeted, or each named organization suffered a breach. [1][2]
- 3 **UNRESOLVED** — The full customer list, ultimate operators, total infection population, legal responsibility throughout the chain and effectiveness of all vendor safeguards remain open. Researchers call for investigation; their conclusions are not judicial findings. [1]

SECTION 03

Figure 01 — The Two-Track Media Trust Boundary

The audience sees an approved message travelling from brand to agency to placement. In parallel, the delivery infrastructure exchanges data and may execute third-party code. Artefact91's conceptual figure distinguishes public-facing communication from its less visible technical delivery path.



SECTION 04

Artefact91 Analysis — Four Strategic Fault Lines

- 1 **REPUTATION WITHOUT DIRECT CONTROL** — Brands may be associated with infrastructure they do not administer. Yet executives will be asked which partners were screened and how quickly the organization responded. Accountability should be calibrated to contractual control and evidence of knowledge, not mere presence on the same advertising network.
- 2 **APPROVAL IS NOT AN AUDIT** — An approved-vendor listing may document admission rather than continuing verification of ownership, subcontracting, dynamic code or downstream conduct. Media plans should distinguish vendor acceptance from recent, independently substantiated security assurance.
- 3 **MEDIA BUYING MEETS SOURCE PROTECTION** — For newsrooms, political offices, advocacy groups and high-risk institutions, covert profiling can threaten confidential sources, stakeholder participation and staff safety. Agencies serving these audiences should no longer separate media delivery assurance from crisis preparedness.
- 4 **THE SILENCE PROBLEM** — A hasty public statement can exaggerate harm; an exclusively technical response can leave stakeholders without a credible explanation. Leaders need a disciplined way to distinguish what is known, what is being tested, who is accountable and when the next update will arrive.

SECTION 05

The Artefact91 Media Trust Boundary Test

Seven decision gates translate the evidence into proposed safeguards for communications teams, publishers and institutional campaigns. These are governance recommendations—not guarantees against sophisticated attacks.

GATE 1 MAP	Identify the principal, agency, demand-side platform, exchanges, supply-side platforms, publishers, verification services and tag operators. Require actual organizations, roles and subcontractors rather than a generic platform description.
GATE 2 VERIFY	Request counterparty screening, ownership information, security reviews, independent audits, incident history, subcontractor disclosure and enforceable suspension rights. A platform admission badge is not conclusive due diligence.
GATE 3 MINIMIZE	Document the need for each targeting variable, identity linkage and sensitive audience segment. Test whether aggregate, contextual or less personalized buying can meet the communications objective.
GATE 4 INSPECT	Assess visible creative separately from tags, scripts, software kits and remotely loaded content. Ask how controls handle advertising code that changes by IP address, device, geography or time.
GATE 5 MONITOR	Require suitable anomaly detection, repeated scanning, vendor-list change alerts, technical incident indicators and an escalation path. A placement screenshot does not establish what code executed.
GATE 6 RESPOND	Pre-agree the accountable executive, cybersecurity lead, legal reviewer, client contact, evidence-preservation plan and ability to disable or change a buying route promptly.
GATE 7 PROVE	Retain documentation of buying routes, partner checks, testing outcomes, incident decisions and remediation. Contractual audit and exit rights must be practical enough to use during an incident.

THE BOUNDARY RULE

Consider pausing or redesigning a high-risk activation if critical intermediaries cannot be identified, material executable elements cannot be assessed or serious warnings cannot be investigated with credible evidence. This is Artefact91's proposed risk standard, not a legal conclusion.

SECTION 06

Responsibility Without Blame by Association

BRAND / INSTITUTION	Set the acceptable risk profile, ask for transparent supply paths, establish audit rights, and approve escalation; do not make security assurances without technical evidence.
AGENCY / MEDIA BUYER	Map the supply chain, preserve procurement records, challenge supplier assurances and tell clients about material emerging risks.
PLATFORM / AD-TECH PARTNER	Demonstrate screening, malicious-code controls, subcontractor oversight and enforceable suspensions. Request proof, not just a policy webpage.
PUBLISHER / HOST	Review ad-serving dependencies, keep relevant delivery records and coordinate with exchanges and independent specialists when a suspect ad is reported.
SECURITY / LEGAL	Determine technical impact, manage containment and forensic analysis, assess notification responsibilities, and distinguish a supplier allegation from an established incident.

SECTION 07

The First 72 Hours — A Decision Protocol

0–6 HOURS / CLASSIFY — Distinguish a suspect vendor association, unverified ad exposure, confirmed malicious creative and suspected device compromise. Convene the media owner, IT/security, legal and communications leads. Preserve relevant delivery identifiers, creative tags, timestamps, supplier correspondence and logs where feasible. Where there is an active threat, security containment takes priority.

6–24 HOURS / VERIFY — Obtain technical trace data from platforms and publishers. Do not collapse impressions, targeting and infection into one number. Determine whether media buying, website hosting, staff browsing or confidential source protection creates the exposure. Engage qualified incident responders where the risk warrants it and assess contractual and regulatory duties.

24–72 HOURS / ACCOUNT — Brief affected clients and organizational leaders on verified facts, uncertainty and protective actions. Set an update time. Record why an activation was stopped, continued or changed. Use an evidence register reviewed by security specialists to approve public claims.

Illustrative Holding Language

“We are reviewing whether any relevant media or technology partners were connected to the reported activity. The researchers' findings do not themselves establish that our organization, systems or audiences were compromised. We are working with appropriate providers and specialists to verify any material exposure and will update relevant stakeholders if confirmed information warrants it.”

SECTION 08

Decision Intelligence — Prioritize the Response

Assess three separate dimensions: **(A)** proximity to the implicated media-delivery pathway, **(B)** potential consequences, including sensitivity of staff or stakeholders, and **(C)** confidence in evidence. High consequences plus credible proximity call for immediate technical scrutiny even if infection is unconfirmed. Low proximity without concrete indicators calls for monitoring rather than a crisis declaration.

Useful measures include the percentage of media spend with a mapped delivery chain, share of executable tags reviewed independently, age of supplier checks, time required to retrieve impression-level records, warning-to-containment decision time and outstanding high-risk intermediaries. These are assurance metrics, not estimates of spyware prevalence.

Board and Executive Questions

- 1 Can the media lead and security lead jointly explain the delivery route for our highest-risk campaigns?
- 2 Which intermediaries remain unknown to us, and do we have enforceable rights to identify and exclude them?
- 3 Could our approved message appear while a separate vendor executes code we cannot inspect?
- 4 If a reporter asks whether a suspect ad reached our environment, who can distinguish exposure, targeting and compromise?
- 5 What verified evidence would cause us to pause activity, notify a client or issue a public statement?

ARTEFACT91 BRIEFING

What to Watch Next

NEAR TERM — New technical indicators, vendor audit changes, independent evidence of takedowns and regulatory or cybersecurity-agency responses. Citizen Lab says a related investigation into mass deanonymization is forthcoming; this is a research lead, not proof of additional incidents. [1]

MEDIUM TERM — Whether public-interest campaigns adopt media-security requirements; whether publishers can insist on executable-code transparency; and whether assurance standards move beyond vendor self-attestation.

Why This Qualifies as a Special Briefing

Publication outside Artefact91's normal twice-monthly cadence is justified by a newly documented attack channel and immediate implications for media procurement, newsroom protection and crisis readiness. The distinctive issue is trusted advertising infrastructure used for surveillance and exploit delivery—not a routine AI-content, election-messaging or executive-impersonation development. This analysis does not duplicate earlier Artefact91 briefings.

ARTEFACT91 BRIEFING

Sources & Research Note

[1] Marczak, B., Christl, W., Scott-Railton, J., Perry, A., Anstis, S., & Deibert, R. (2026, October 15). *Exposing AdFaction: How Discover-Tech Turned Targeted Ads Into a Pipeline for Zero-Click Mercenary Spyware*. Citizen Lab Report 195, University of Toronto. <https://citizenlab.ca/research/how-discover-tech-turned-targeted-ads-into-a-pipeline-for-spyware/>

[2] Confiat Threat Intelligence. (2026, October 15). *ADINT: Weaponized Ad Tech for Zero-Click Spyware and Mass Identity Vacuuming*. Technical investigation; telemetry analyzed covers April 1–September 10, 2026. <https://blog.confiant.com/p/adint-weaponized-ad-tech-for-zero-click-spyware-and-mass-identity-vacuuming>

[3] Satter, R. (2026, October 15). *Israeli firm used online ads to run vast spy operation, researchers say*. Reuters. Includes platform responses and attempts to obtain comment from the investigated firms. <https://www.reuters.com/legal/litigation/israeli-firm-used-online-ads-run-vast-spy-operation-researchers-say-2026-10-15/>

Editorial Method

This briefing synthesizes two concurrent primary technical investigations and separately reported platform responses. Artefact91 did not independently re-analyze malware specimens or network logs. The Media Trust Boundary Test and operational decision rules are original professional recommendations, not empirical findings by investigators. All findings reflect reporting available on October 15, 2026.