CONFIANT

Demand Quality Report

Q4 2019

Introduction

Confiant's **Demand Quality Report** is a quarterly look into the quality of demand in digital advertising. Using a sample of 350 billion impressions monitored in real time, Confiant is able to answer fundamental questions about the state of ad quality in the industry at large.

Digital advertising delivers significant value to publishers but introduces myriad risks related to security and user experience. **Malicious**, **In-Banner Video**, and **Low Quality ads** diminish the value of demand and drive user adoption of ad blockers. However, few if any systematic studies have been conducted 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 has historically been challenging to estimate impact without client-side instrumentation in place on a large and diverse set of publishers. The Demand Quality Report, which leverages Confiant's position as the vendor of choice for real-time creative verification, aims to change that.

In October 2018, Confiant released the industry's first benchmark report. This report, the seventh in the series, covers the entirety of 2019, with a particular focus on Q4.

Definitions

Malicious ad A creative that includes (usually obfuscated) Javascript that spawns a forced redirect or loads a secondary, or tertiary, payload for similar malicious purposes. Most malicious creatives exist for the purpose of forcing users to interact with phishing scams, but some perform cryptojacking or infect the user's device to propagate botnets and other nefarious activities.

In-Banner Video (IBV) ad The practice of serving video ads in banner placements without the publisher's consent, and often without the advertiser's consent either. In these cases, a video ad unit is loaded within a banner placement as a display unit, instead of playing within a media player.

Low Quality ad Creative violations across a range of different quality specifications selected by the publisher. The dimensions include audio/video related violations, creatives probing for user's geolocation, the network load of the ad, and much more.

Methodology

To compile the research contained in this report, Confiant analyzed a **normalized sample of more than 355 billion programmatic advertising impressions** from January 1 to December 31, 2019.

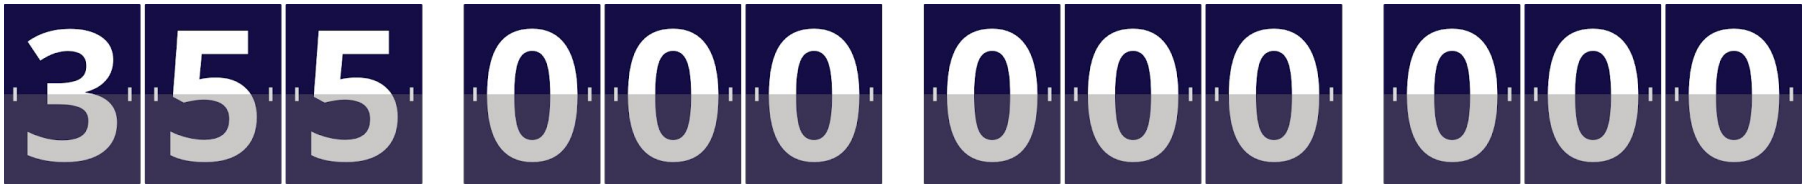

The data was captured by Confiant's **real-time creative verification solution**, which allows us to **measure ad quality on real impressions for real users** across devices and channels.

84 Billion Problematic Impressions Industry-wide in 2019

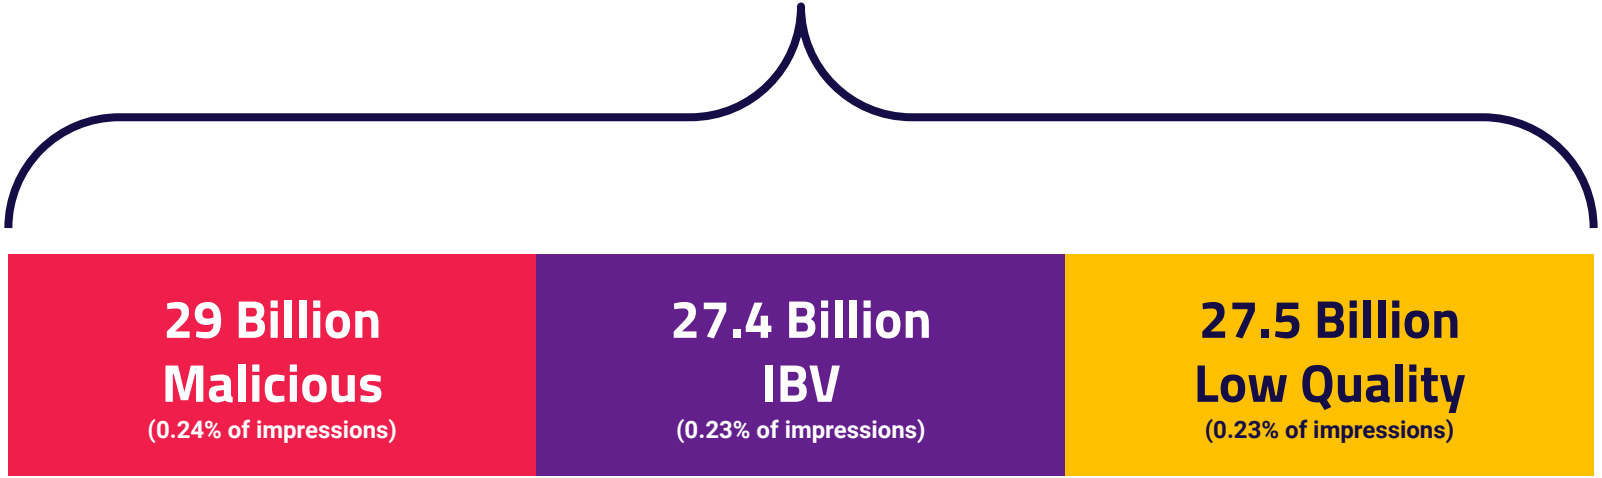

0.70% of impressions in our sample were marred by a **dangerous** or **disruptive** ad. When taken across the enormous scope of programmatic advertising, this equates to **84 billion problematic impressions**.

How did the industry do in 2019?

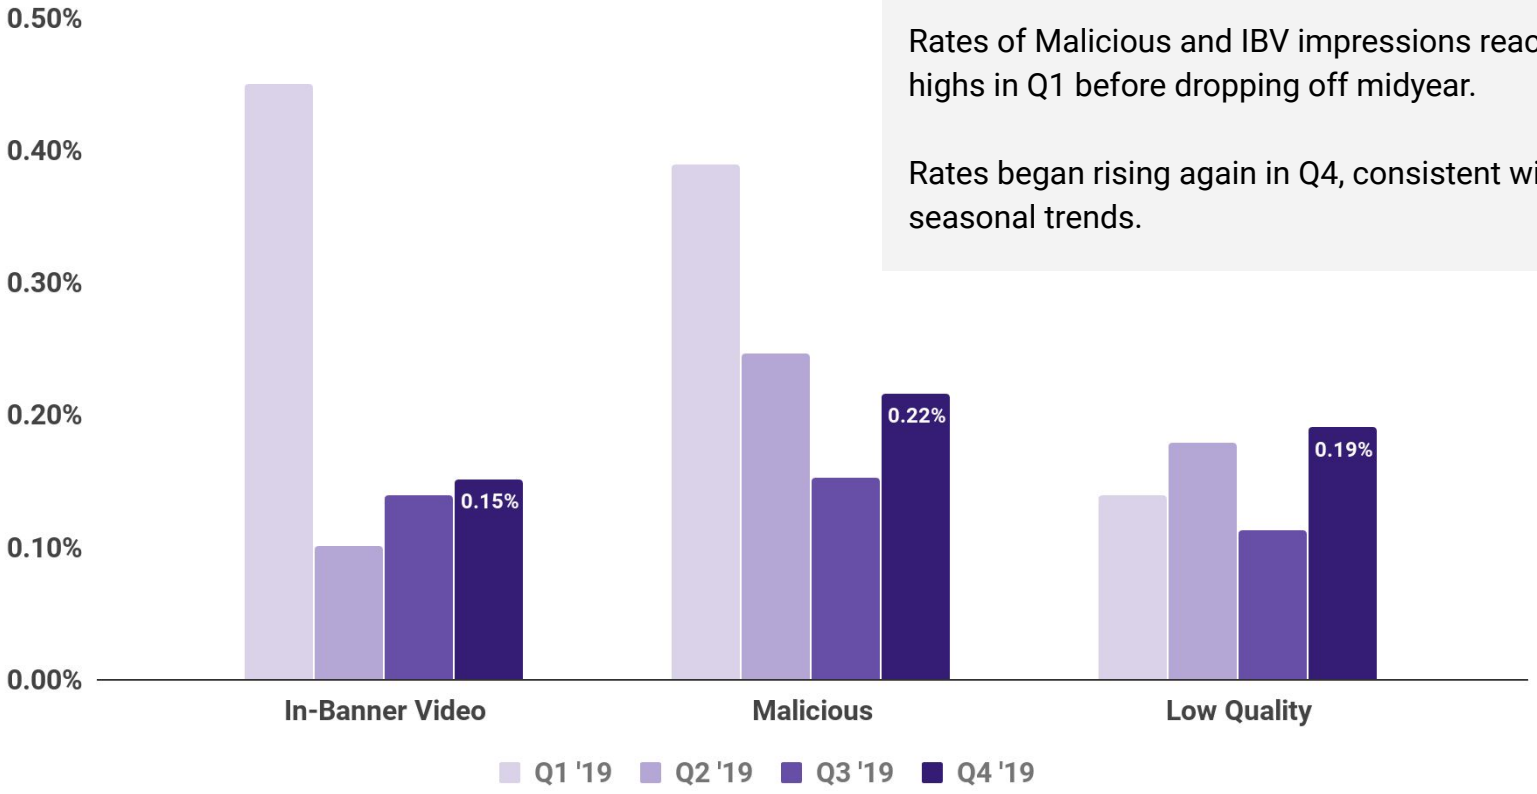

Rates of Malicious and IBV impressions reached highs in Q1 before dropping off midyear.

Rates began rising again in Q4, consistent with seasonal trends.

Issue frequency by user agent: 2019

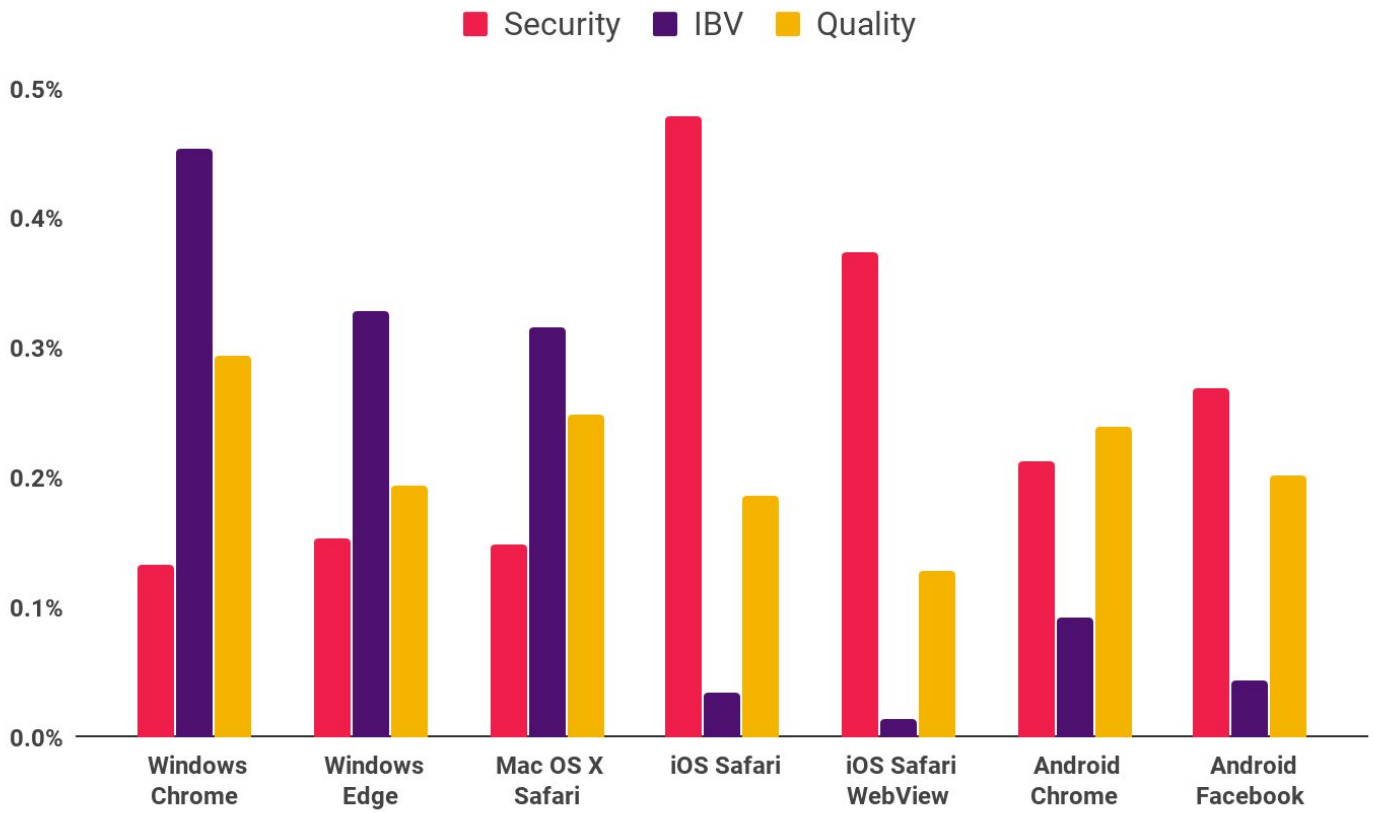

The frequency of ad quality issues varies considerably across browsers.

For 2019 as a whole, **iOS Safari had the highest rate of malicious ads** by a significant margin.

Conversely, IBV was far more common on desktop browsers such as Chrome and Edge for Windows.

In 2019, **1 in every 150** impressions was **dangerous** or **highly disruptive** to the end user.

Q4 US SSP Rankings

In Q4, Confiant tracked impressions from over 75 SSPs. However, **over 80% of impressions originated from just 13 providers**[1] commonly used by publishers. These providers are noted in the charts that follow using a coding system that carries over from one quarter to the next.

To qualify for inclusion, a provider had to be the source of **at least 1 billion impressions** across our cross-section of publishers.

We identify Google Ad Exchange within these rankings. As the operator of the largest exchange, Google has access to data and resources beyond what's available to other exchanges, which one could reasonably expect to translate into higher efficacy when it comes to catching issues.

Our data confirms this assumption, with Google Ad Exchange consistently placing among the top performers.

[1] Google AdX, Rubicon Project, OpenX, Xandr, Verizon Media, Index Exchange, Pubmatic, EMX, Sonobi, TripleLift, District M, 33Across, and Sovrn

Malicious Impressions by SSP Q4 ‘19

We continued to see huge disparities in performance among the top SSPs: **the worst, SSP-I, delivered malicious ads at 33x the rate of the best**.

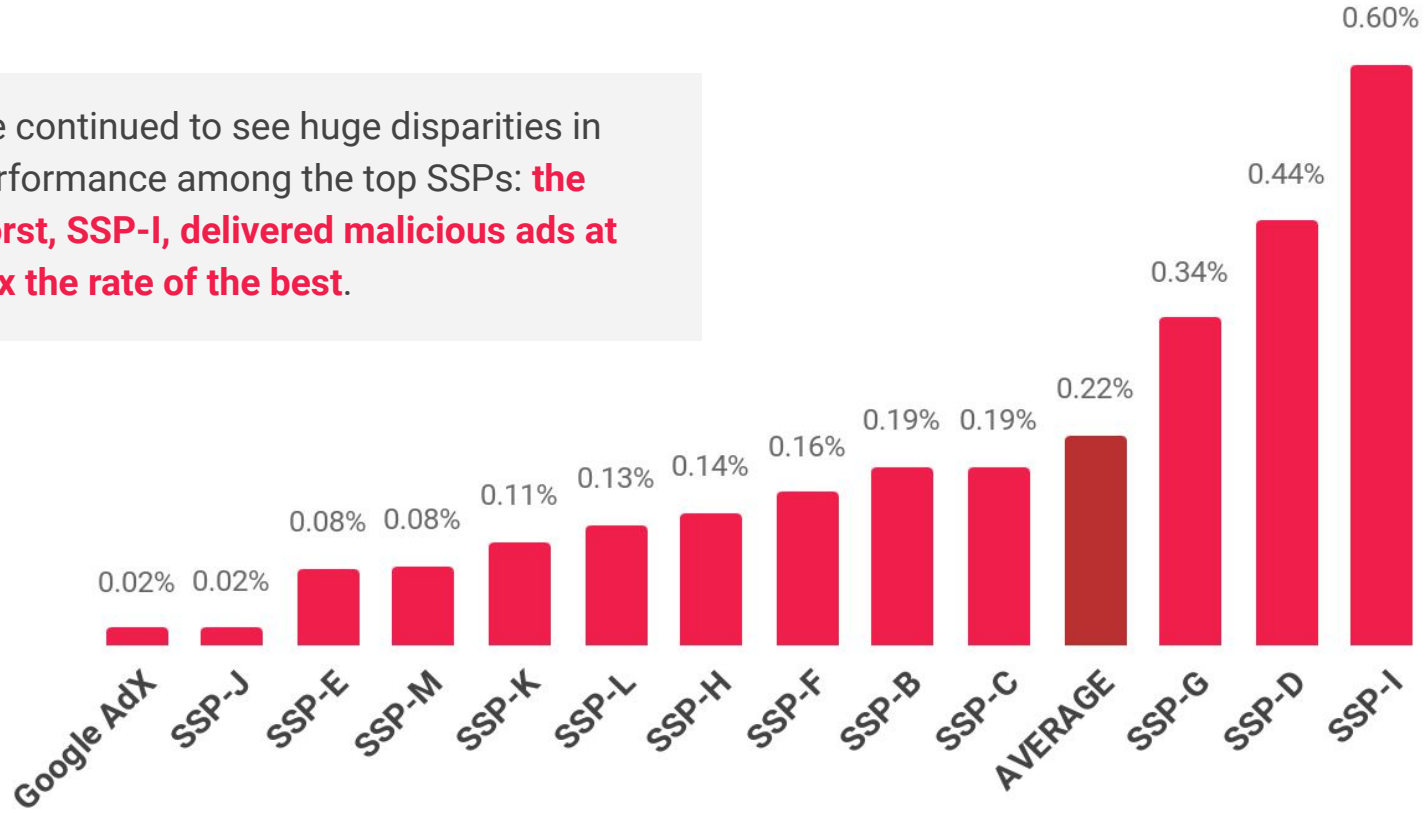

Malicious Ads by SSP: High-Risk Ad Platforms

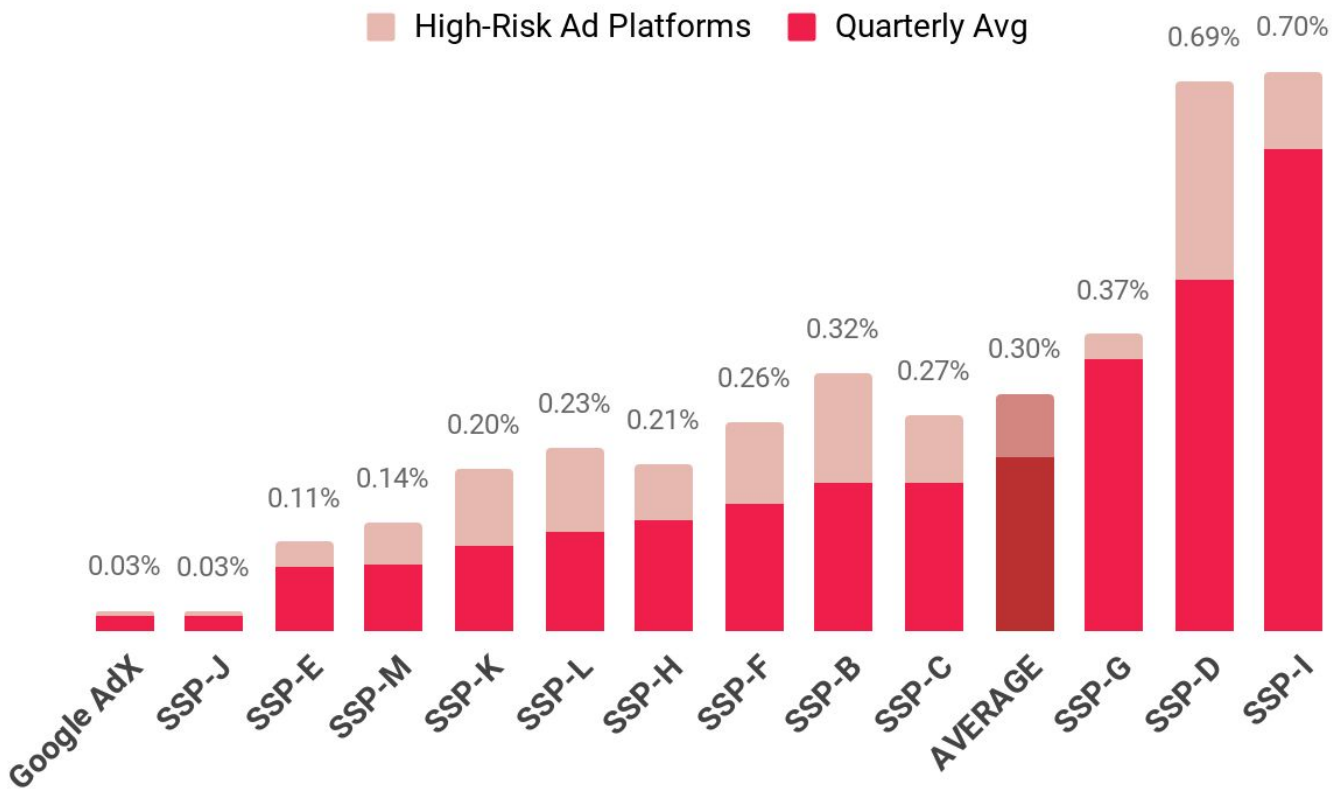

High Risk Ad Platforms (HRAPs) are ad platforms that consistently deliver high-risk creatives.

While HRAPs offer some safe demand, their high rates of malicious ads make it a wise decision for most publishers to block them.

This chart shows how the inclusion of impressions from HRAPs changes the overall ranking and illustrates the **strong correlation between malicious rates and HRAP activity.**

Malicious Ads by SSP: Q3 to Q4

Four SSPs showed big increases in Malicious ad rates from Q3 to Q4, while Google returned to its historically strong performance.

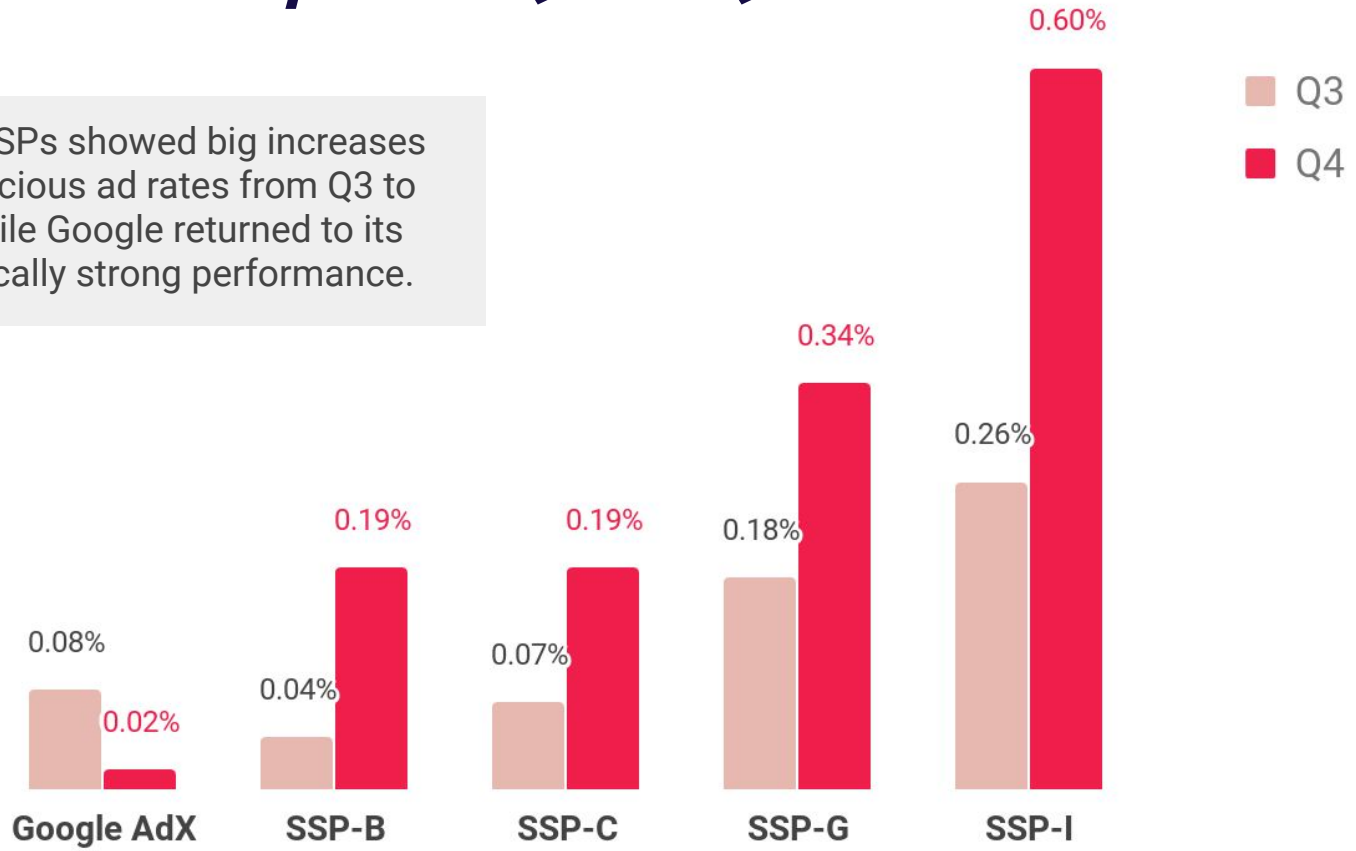

Daily Maximum Malicious Rate by SSP

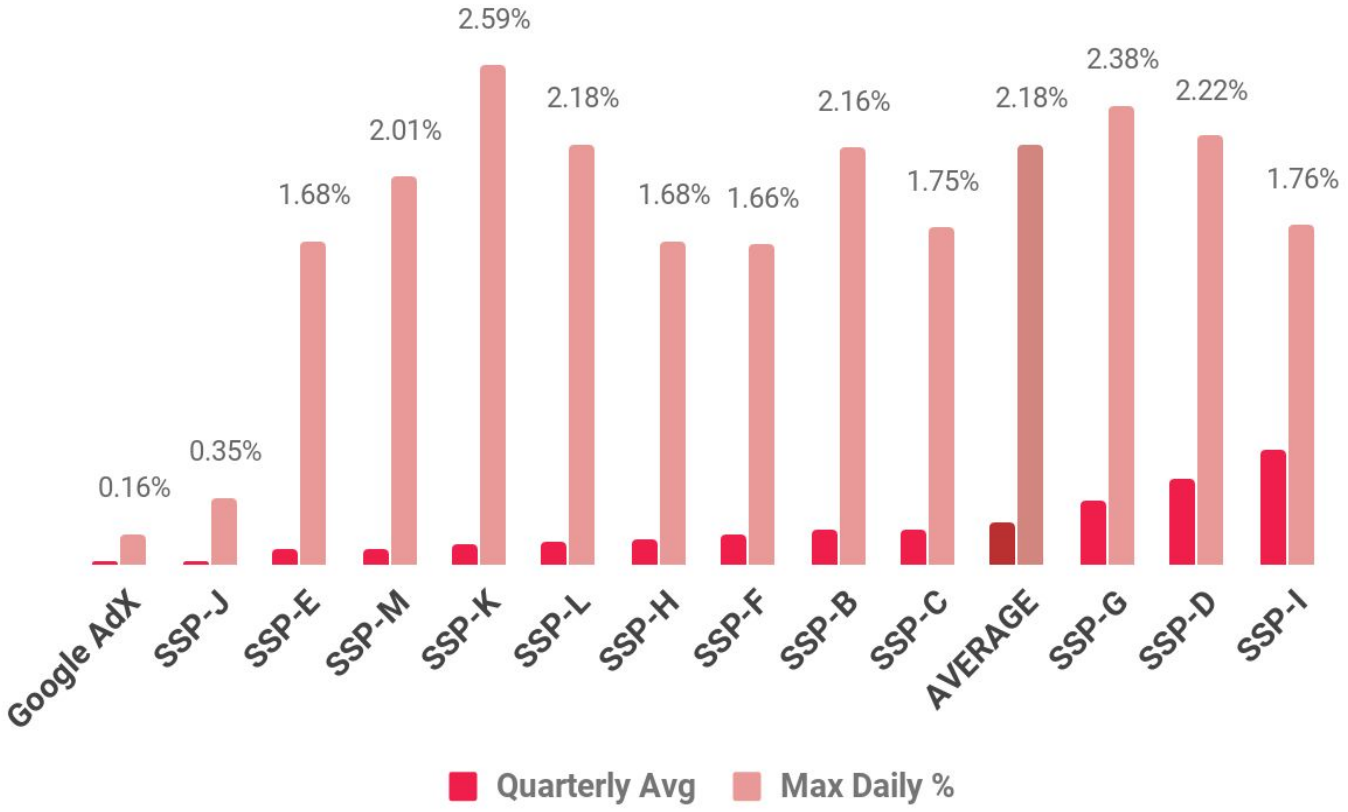

Quarterly averages can mask significant variation in day-to-day performance, so it's important to measure the **upper bound of the Malicious ad rate** for each SSP to get a complete sense of performance.

When under sustained attack, **SSPs had days when 2 of every 100 impressions were malicious**, putting publisher relationships at considerable risk.

Avg Malware Attack Response Time (in days)

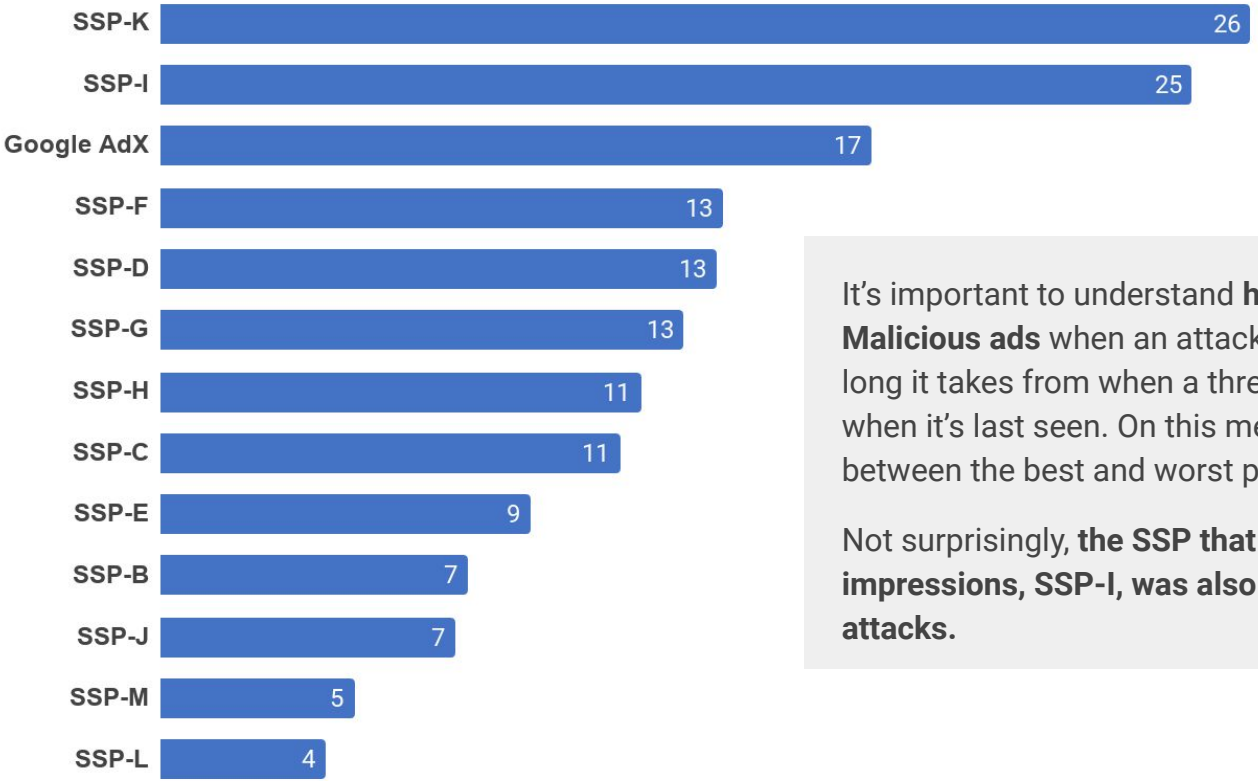

It's important to understand **how quickly an SSP responds to Malicious ads** when an attack is underway. We measure how long it takes from when a threat first appears on an SSP to when it's last seen. On this measure, we see huge disparities between the best and worst performers.

Not surprisingly, **the SSP that had the highest rate of Malicious impressions, SSP-I, was also among the slowest to respond to attacks.**

Malicious ads: SSP Performance in 2019

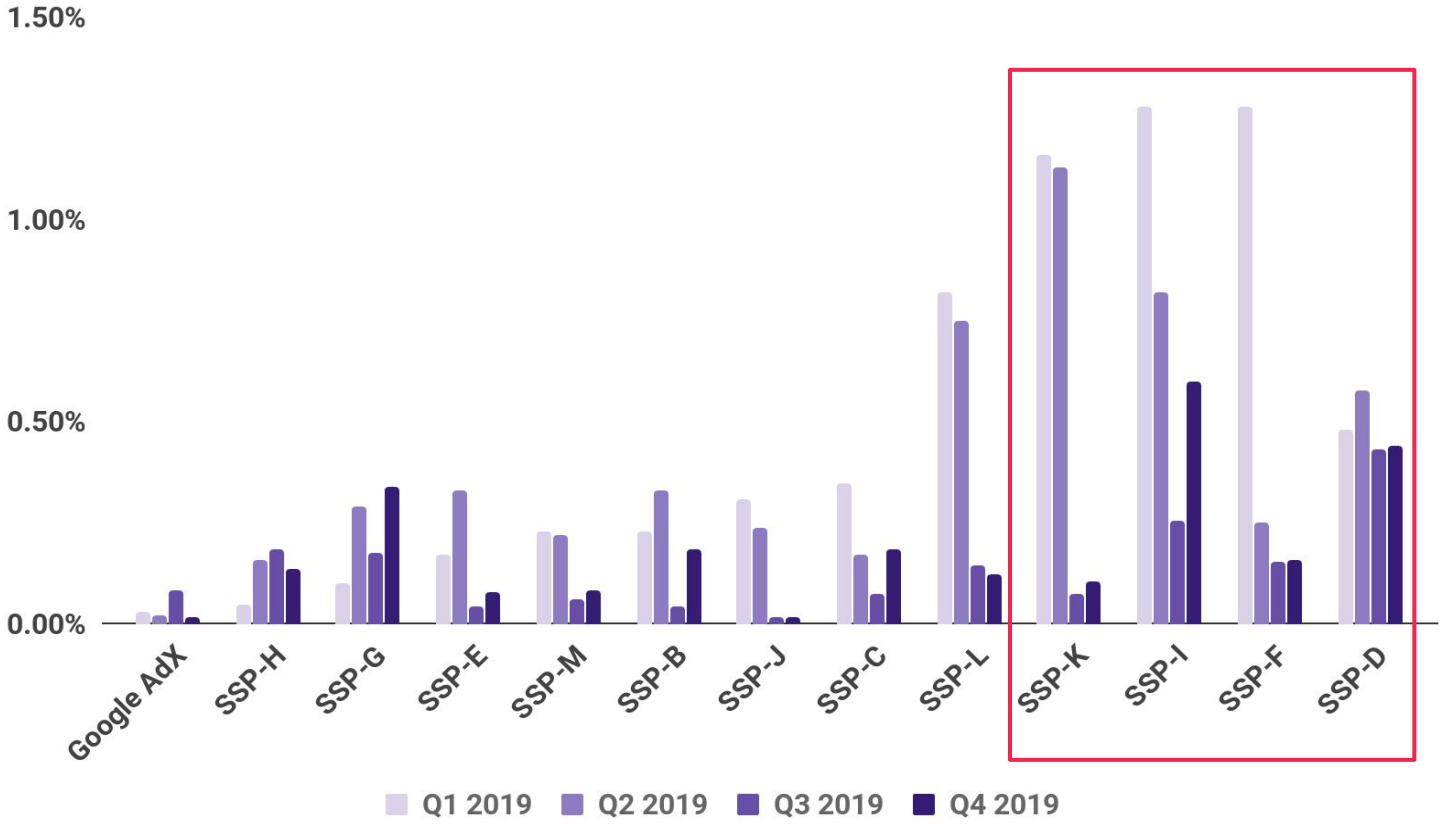

When it comes to Malicious ads, SSPs show meaningful variance in performance from quarter to quarter, with even Google showing occasional spikes in activity.

SSP-K, SSP-I, and SSP-F showed the largest swings in performance. Conversely, **SSP-D was a perennially poor performer**, calling their partner vetting and management practices into question.

Low Quality and In-Banner Video Ads Q4 ‘19

Low Quality ads are based on a diverse set of rules that publishers can elect to activate on the Confiant platform. Examples include **autoplay audio, heavy ads**, **pop-ups**, and **In-Banner Video**. These rules correspond to ad behaviors that have one feature in common: they disrupt or impair the user experience.

SSP-D is now the lone significant source of unwanted In-Banner Video, making them an obvious choice to disable for quality-focused publishers.

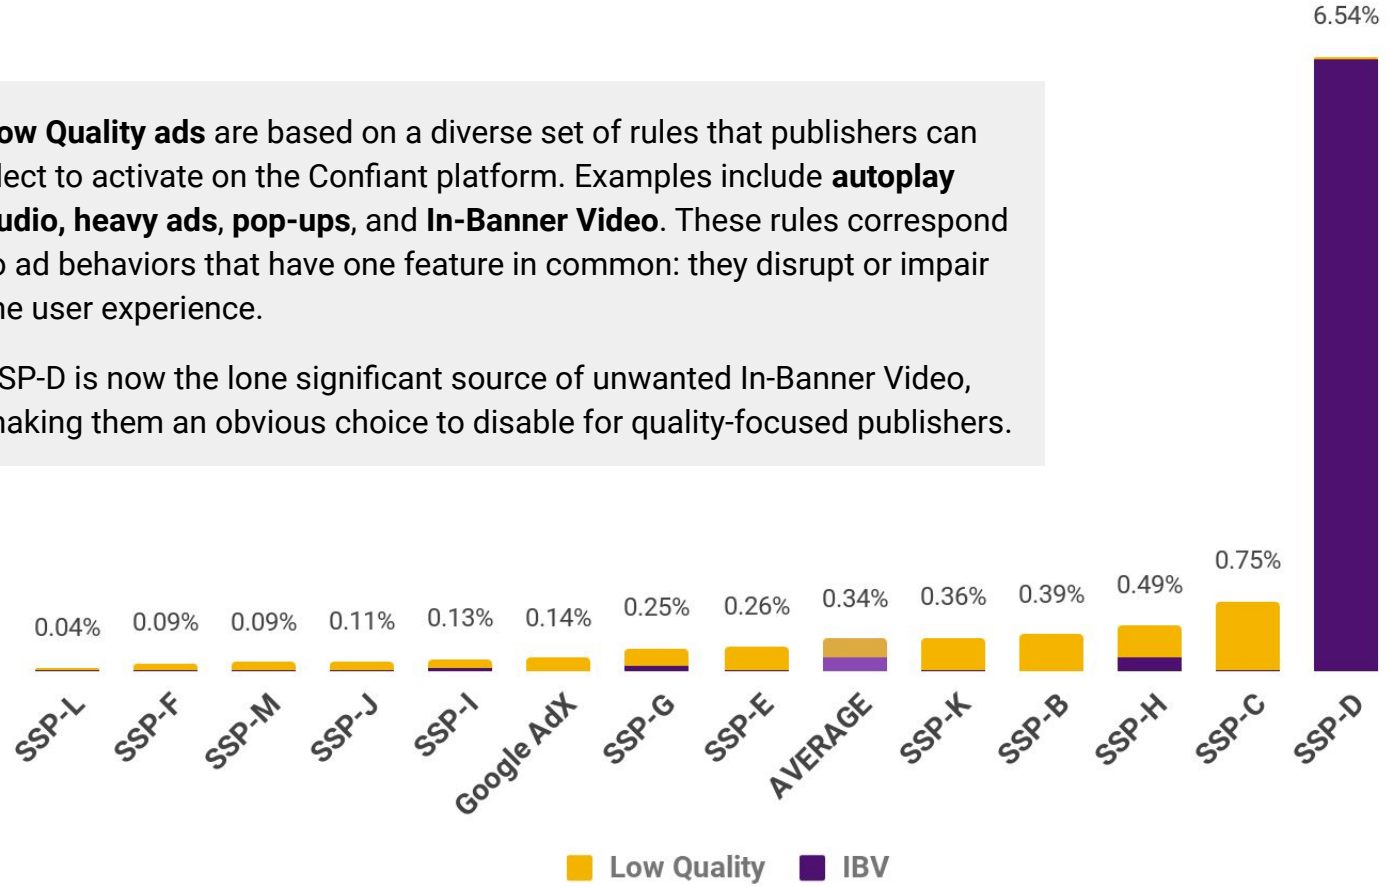

What quality issues do publishers care most about?

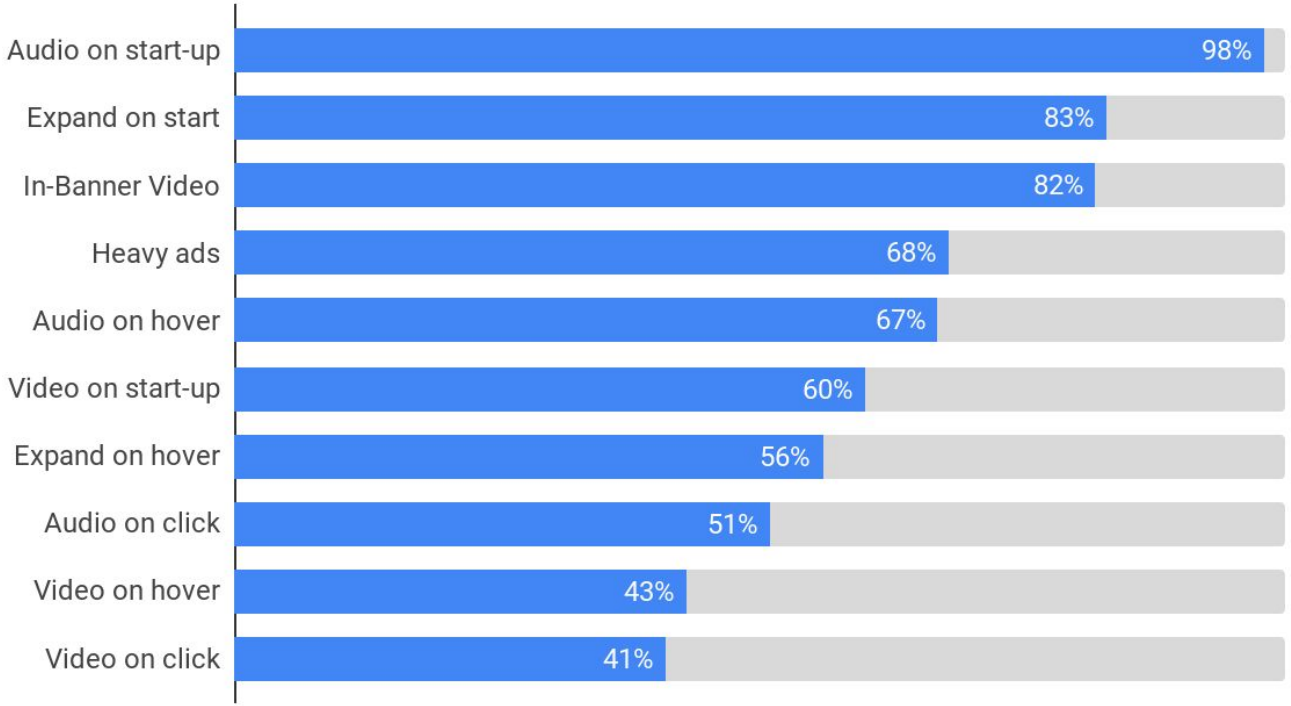

Confiant publishers can choose whether to activate monitoring for quality issues in accordance with their business needs and the expectations of their audience. This chart summarizes the rate at which various rules were activated by our publishers for blocking across all impressions monitored by Confiant in Q4.

Unsurprisingly, **blocking activation rates tend to be higher for automatic creative behaviors** (e.g. Audio on start-up) than those requiring user action (e.g. Audio on click).

In Q4, **50%** of low-quality and IBV impressions came from just **2 SSPs**

The collapse of In-Banner Video

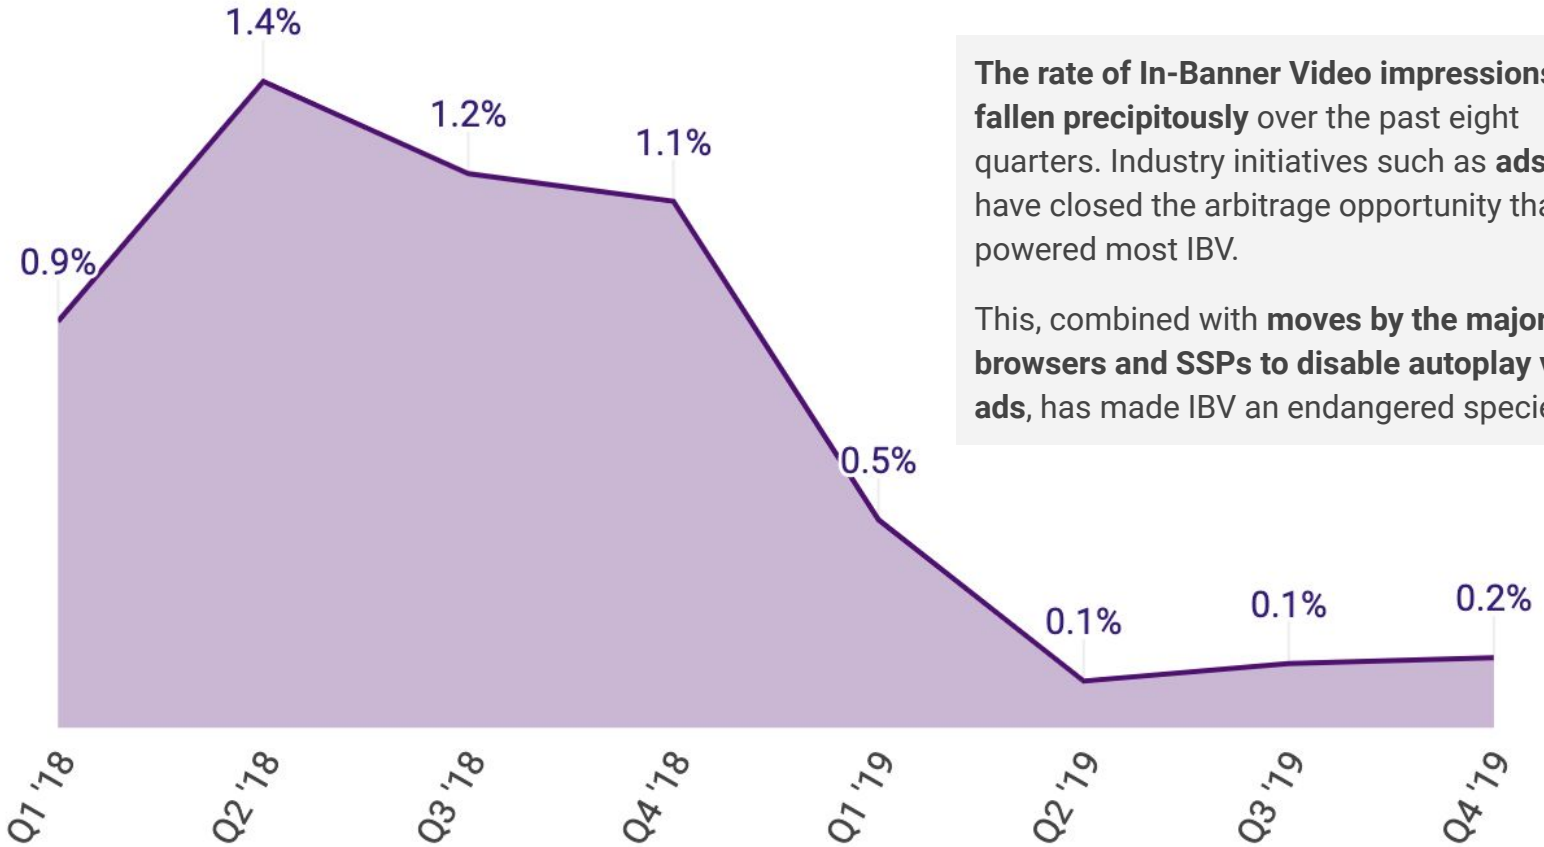

The rate of In-Banner Video impressions has fallen precipitously over the past eight quarters. Industry initiatives such as **ads.txt** have closed the arbitrage opportunity that powered most IBV.

This, combined with **moves by the major browsers and SSPs to disable autoplay video ads**, has made IBV an endangered species.

Q4 Rates by Country

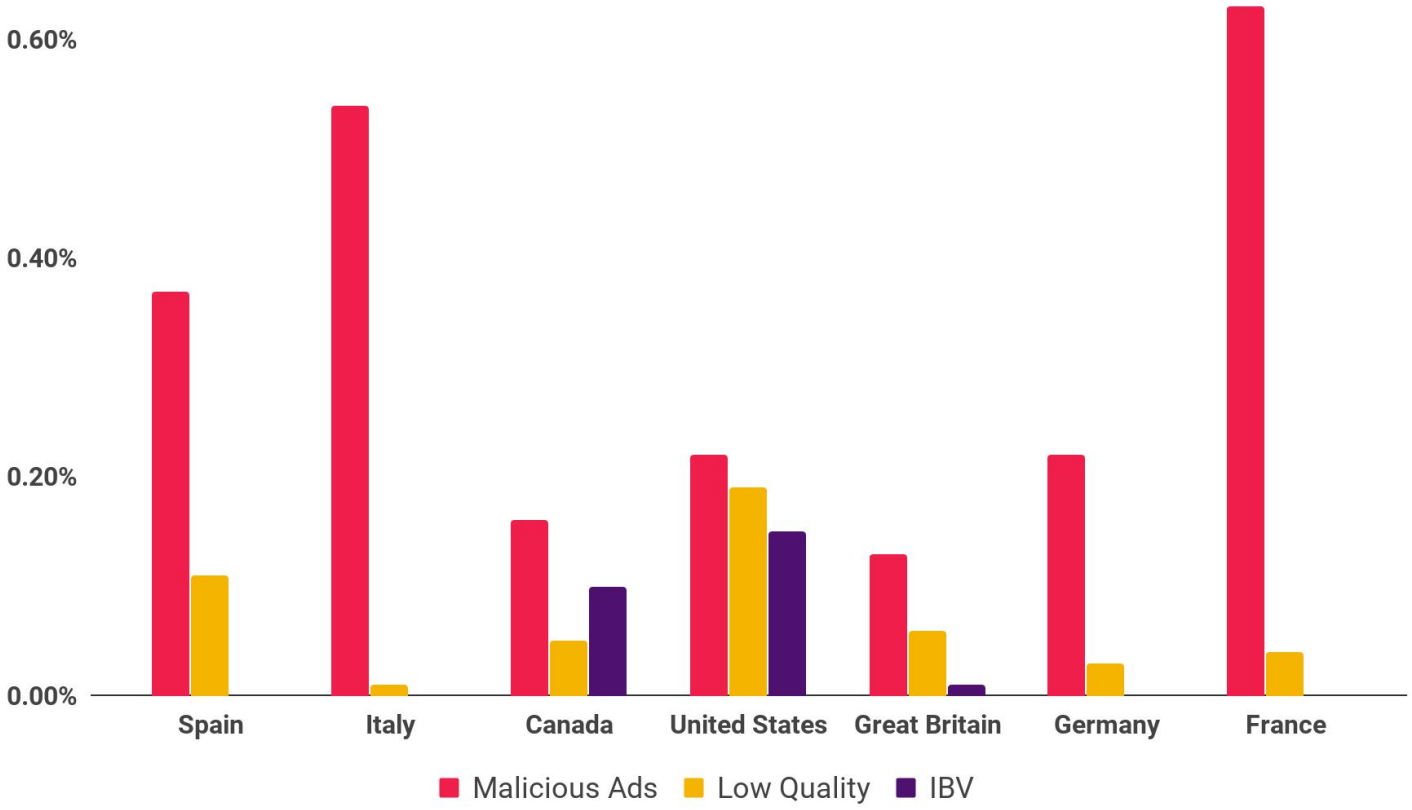

As in past quarters, **European markets saw far higher rates of Malicious ads than the U.S.**, but a lower rate on other issues.

The variety of rates by country exemplifies how malvertisers continually shift their campaigns and targets to remain under the radar.

Major Threat Groups Active in Q4

Threat Activity

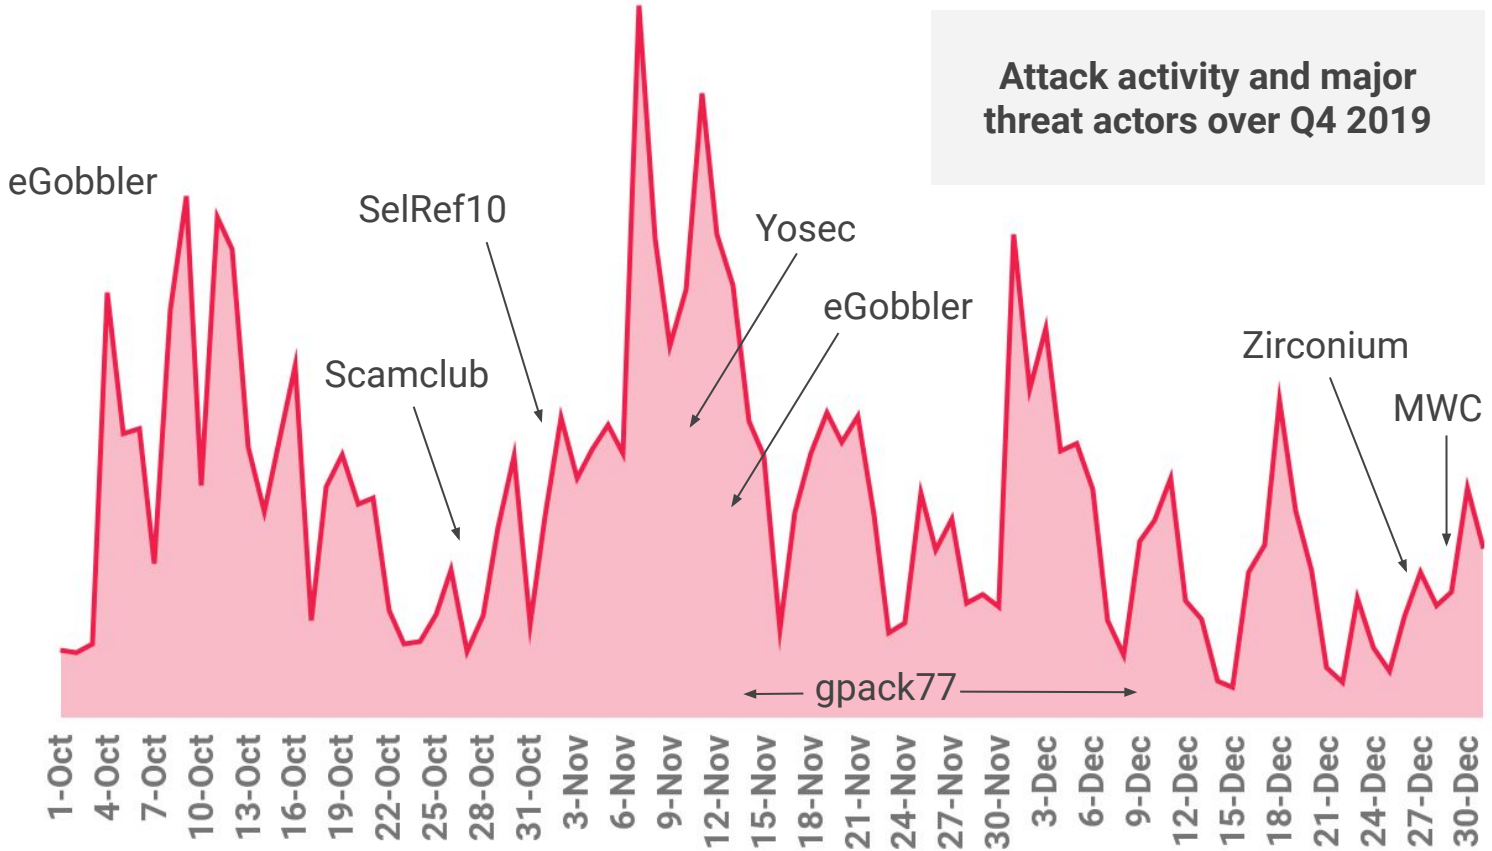

eGobbler

Peak activity: October

Notable characteristics: This Asia-based attack group has a history of exploiting obscure browser bugs to bypass built-in browser protections against pop-ups and forced redirects. After Confiant discovered a previous vulnerability in early 2019 and worked with the Chrome team to shut it down, eGobbler introduced a Webkit exploit.

Early in October, we saw a brief surge of activity, followed by a notable campaign running on Oath that had utilized a Polyglot evasion tactic - which means the payload code was hidden conveniently in an image file.

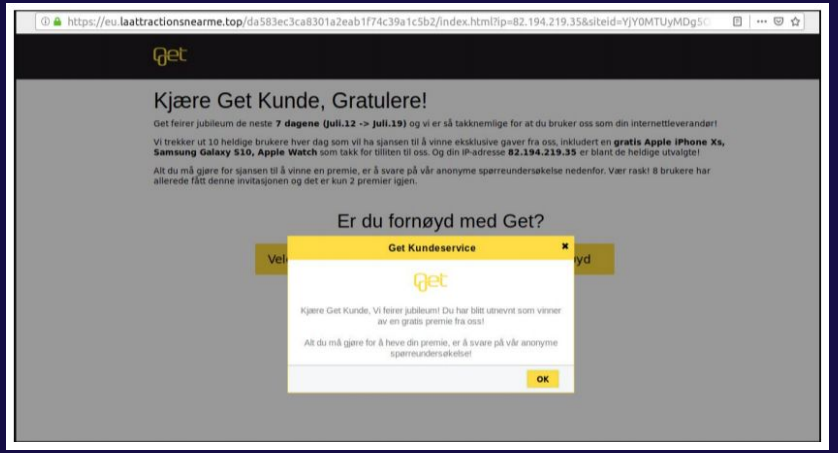

Scamclub

Peak activity: October, December

Notable characteristics: Scamclub stands apart from their malvertising peers in their approach toward evasion. Whereas most high-profile malvertisers choose to hide behind carefully crafted fingerprinting and targeting, Scamclub relies on cranking out dozens (or hundreds) of creatives daily with subtle variations in very rudimentary obfuscation.

This bombardment tactic is designed to overwhelm platforms and security vendors by creating a flood of dangerous demand that they hope will spill beyond any anti-malvertising gatekeeping.

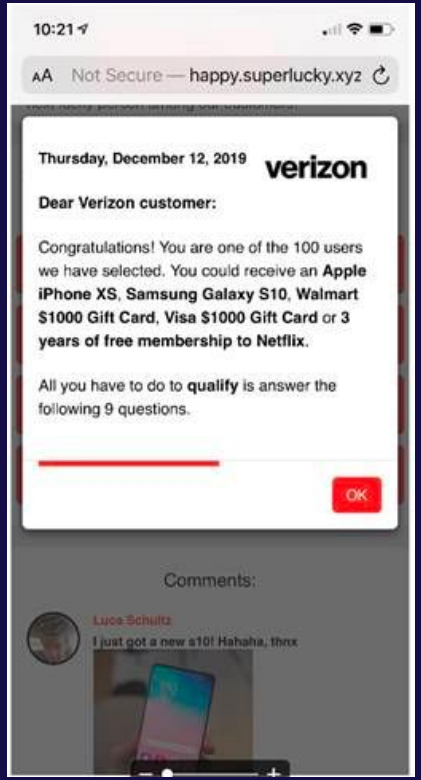

Yosec

Peak activity: October, December

Notable characteristics: One of the new kids on the malvertising block, Yosec had close to 100 unique malicious creatives active as of 10/31, predominantly via AdForm.

Named after their CTA messaging along the lines of “Your Mac Security”, this attacker based in Eastern Europe has been consistently serving up redirects to threatening malware pages.

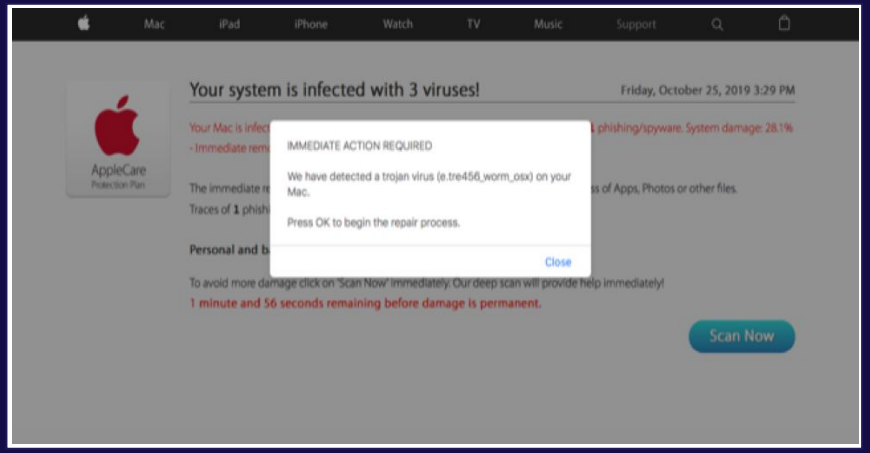

SelfRef10

Peak activity: early November

Notable characteristics: We first called the industry's attention to this malvertising middleman about 6 months ago in our blog post here.

SelfRef10 specializes in forming bi-directional ad tech relationships that empower them both to buy and to sell so that they can play both sides of the coin. They've shown no signs of slowdown in the last half year and continue to run desktop redirect campaigns, often choosing vague domains as delivery vehicles. They've been active through November via Index Exchange on MediaSmart DSP as "ClickFollow Ltd".

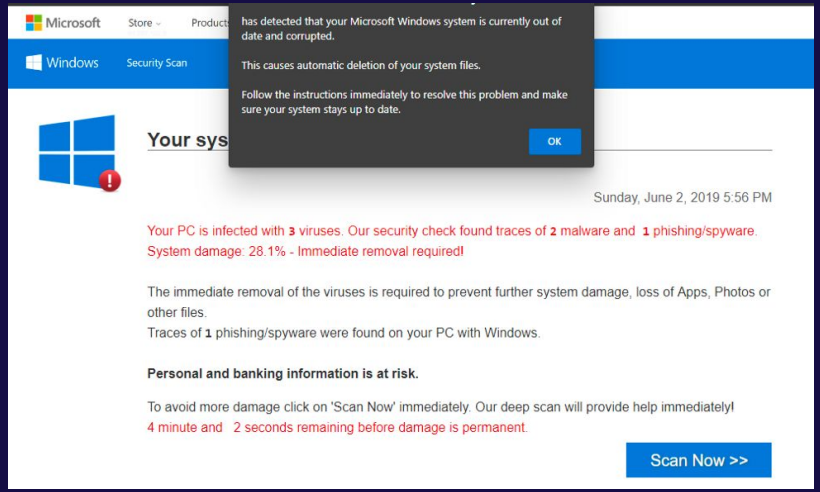

Gpack77

Peak activity: November and December

Notable characteristics: This threat actor based in Latin America is responsible for mobile redirects (Android) that primarily target Europe.

Encountered some major ad serving pivots with a campaign that ran for the first two weeks of December via AppNexus and served up to the tune of 5-8 million reward scam impressions per day.

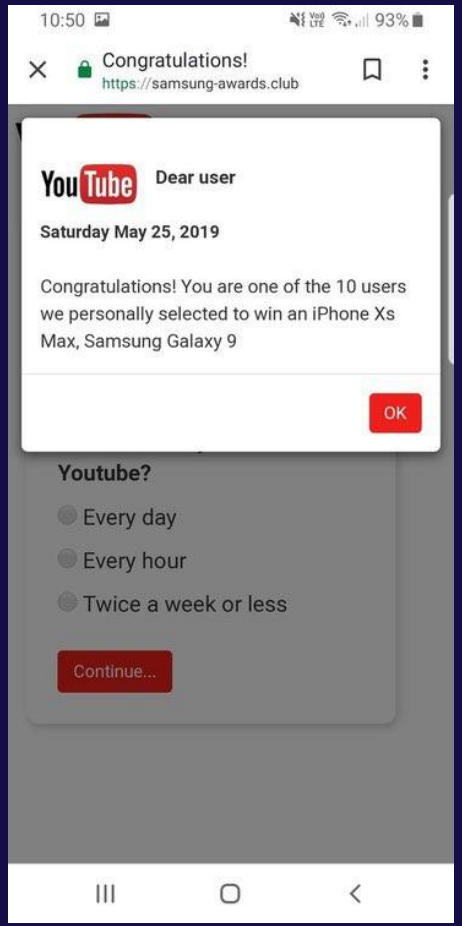

FizzCore

Peak activity: Early December

Notable characteristics: FizzCore is a significant newcomer. An attacker that sits at the increasingly blurred boundary between malvertising and deceptive ads, FizzCore has perfected the art of audit circumvention to exploit the gullibility of aspiring cryptocurrency investors.

Eschewing forced redirects, FizzCore implements techniques to evade ad quality reviews and drive users to cybersecurity scam sites.

Evasion techniques include cloaking (display of fake ad creatives and landing pages to ad quality scanners), reputation and relationship building in the ad ecosystem, and carefully crafted localized campaigns using celebrity endorsement clickbait.

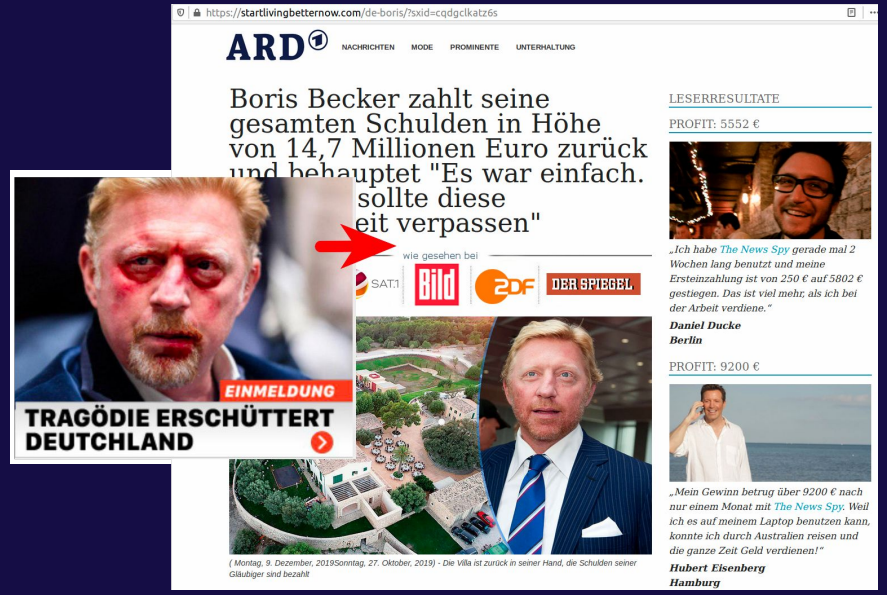

MWC

Peak activity: Late December

Notable characteristics: MWC is the attacker behind some of the ubiquitous giveaway or sweepstakes scams that are often byproducts of force redirect campaigns.

Typically they masquerade as ecommerce brands, going as far as creating an entire digital presence for each of their campaigns by stealing existing brand's ecomm websites and hosting them under typo domains. They have become so adept at "brand identity theft" that they can be indistinguishable from the original except for very minor details.

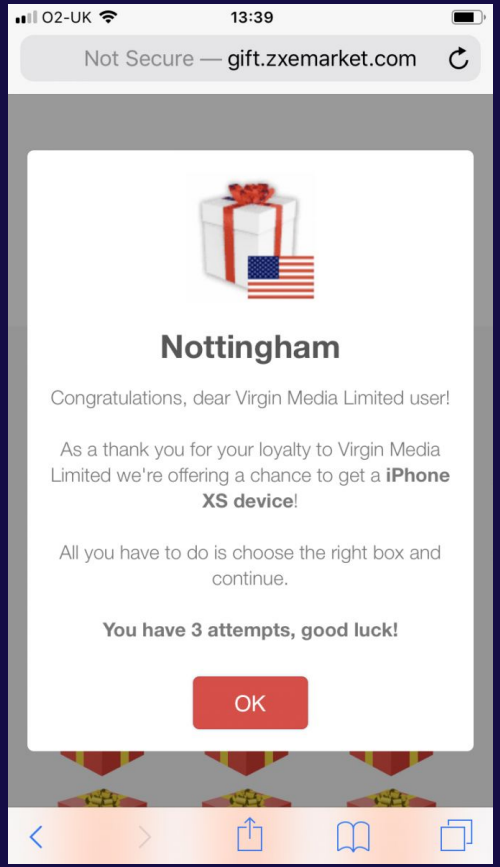

Zirconium

Peak activity: Late December

Notable characteristics: Zirconium runs a very sophisticated malvertising operation that's notable for unique fingerprinting techniques that are carried out in multiple stages. This group, which just two years ago was focused on churning out fake agencies by the handful in order to win seats on buying platforms, has since shifted their approach, but are still running similar tech support focused malvertising campaigns.

The attacker stands out in their choice to target primarily desktop devices and their use of increasingly sophisticated Javascript obfuscation.

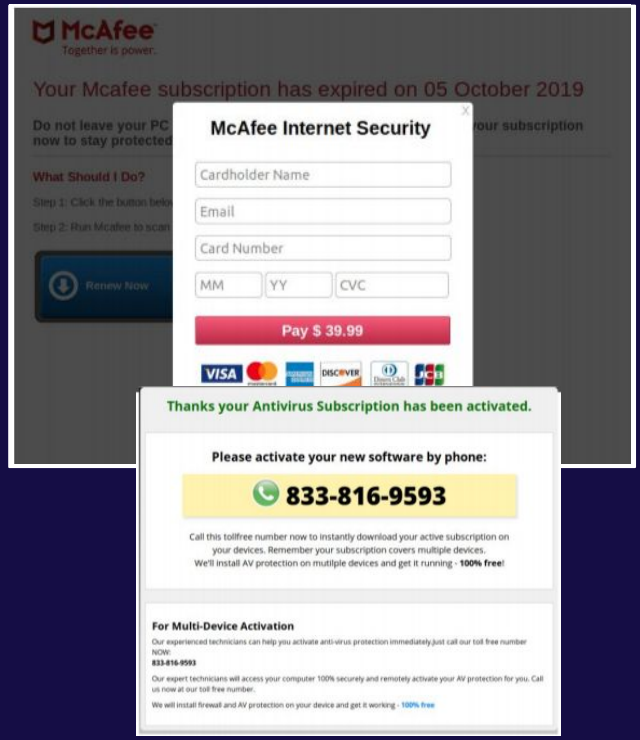

Conclusion

2019

For 2019 as a whole, we detected **serious security or quality issues with 1 in every 150 impressions**.

Safari for iOS had the **highest rate of malicious ads** among top browsers.

2019 was the year that **In-Banner Video** was effectively **knocked out** of the marketplace.

The emergence of **FizzCore** heralds a **new line of attack** beyond redirects: the **use of clickbait to drive users to scam pages.**

Q4

The rate of issues across all categories rose significantly from Q3 to Q4 2019.

The **worst-performing top SSP** was **33x** as likely to deliver a **malicious ad** as the best. Over 60% of malicious impressions came from just 3 SSPs and over 50% of serious quality issues came from just 2.

About Confiant

We believe in making the digital world safe for everyone.

Confiant is a cybersecurity company that protects publishers and platforms from malicious actors and puts the control back in their hands to ensure that ads delivered to users are safe and secure. Our sole purpose is to rid the world of cybercriminals, bad actors, and malware.

Our founders, LD Mangin and Jerome Dangu, teamed up in September 2013 to reinvent how the industry tackled malvertising and low-quality ads. The then-current state of technology was at a data disadvantage against the bad actors that couldn't be surmounted without real innovation. That “never done before” innovation took a year to figure out, and in May 2017 Confiant launched the industry’s first real-time verification and blocking solution, giving publishers actual control of what ads are shown to their users.

Learn More