

# Demand Quality Report for Q3 2019

## Introduction

This report provides insight into the quality of demand in programmatic advertising. Using a sample of 120 billion ad impressions monitored in real time, Confiant is able to answer fundamental questions about the state of ad quality in the industry at large. Programmatic 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. This 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 on the state of programmatic demand. This report, the sixth in the series, covers Q3 2019.

## Definitions

In the context of this research, the terms below are referring to the following 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** - IBV refers to 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.

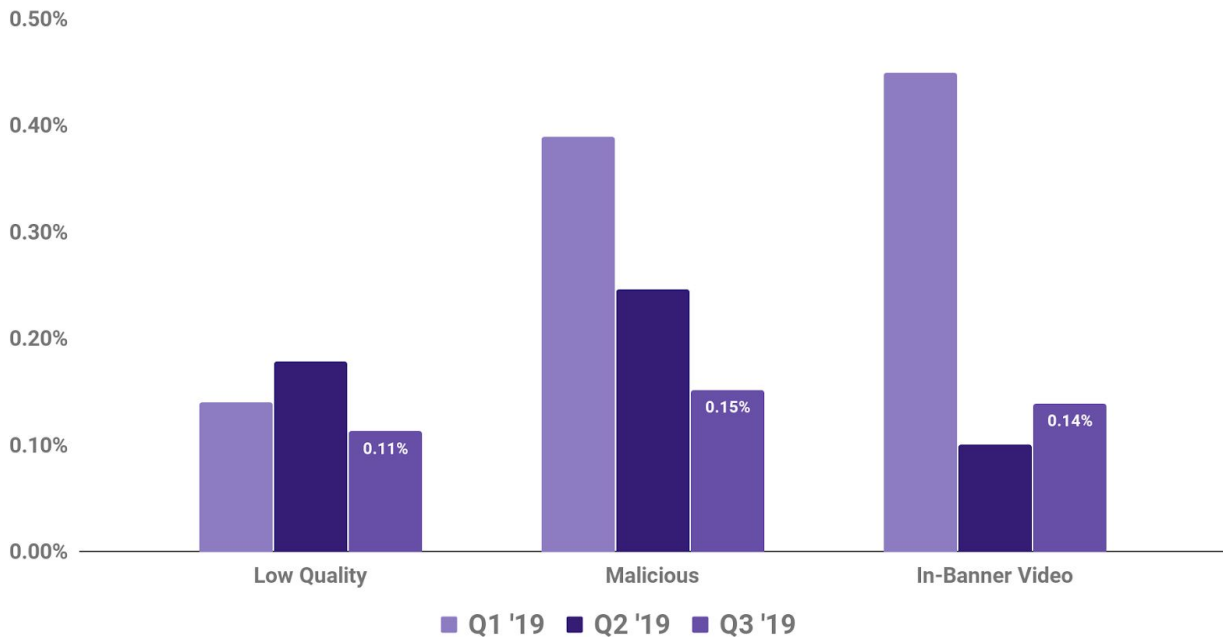
**Low Quality ad** - Ad creative violations across a range of different quality specifications selected by the publisher. The dimensions include audio/video related violations, creatives probing for user's geolocation, the network load of the ad, and much more.

## Methodology

To compile the research contained in this report, Confiant analyzed a normalized sample of more than 120 billion programmatic advertising impressions from July 1 to September 30, 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.

## US Rates by Quarter

The chart below shows Malicious, In-Banner Video, and Low Quality ads as a percentage of total U.S. impressions over the past three quarters. Q3 saw the continuation of a trend that began earlier in 2019: the rate of issues detected has been in decline. On a quarter-over-quarter basis, Malicious ads dropped to 0.15% of total impressions from 0.25% in Q2. Conversely, In-Banner Video ads saw a substantial increase in frequency, coming in at 0.14% of impressions vs. 0.10% in Q2, but remains at low levels by historical standards. Low Quality ads other than In-Banner Video declined to 0.11% of impressions.



Even with these improvements, 1 in every 250 impressions was marred by a serious security or quality issue.

When taken across the enormous scope of programmatic advertising, where up to a trillion ads are served in a given month, such a rate would equate to 4 billion problematic impressions a month.

## Q3 US SSP Rankings

In Q3, Confiant tracked impressions from over 75 SSPs. However, over 80% of impressions originated from just 13 providers commonly used by publishers. To qualify for inclusion in the charts below, a provider had to be the source of at least 1 billion impressions across our cross-section of publishers. We believe that splitting out the data for just top SSPs provides important insights into the performance differences across providers without overwhelming the reader with data. As in past reports, 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 in this and past reports largely confirms this assumption, with Google Ad Exchange consistently placing among the top

performers for each issue type. The section below breaks out performance for these 13 SSPs for Malicious ads and Low Quality ads.


### Malicious Ads

We saw improvement across nearly every provider in Q3 compared to Q2, with one very notable exception: Google Ad Exchange. In previous reports, Google has consistently been the best performer, delivering Malicious ad rates in the 0.02% range. However, Google fell to 7th in the rankings in Q3 and saw their violation rate increase by 320%. Still, they remain a top performer, with their violation rate coming in a 0.08% vs. 0.15% for all impressions monitored by Confiant.

The standout performer in Q3 was SSP-J, who moved into the top spot by reducing their violation rate by an incredible 99% vs. Q3 of last year.


SSP-E and SSP-B, both consistently good performers, moved into second and third place, respectively.

### Malicious Ads



Malicious Ads

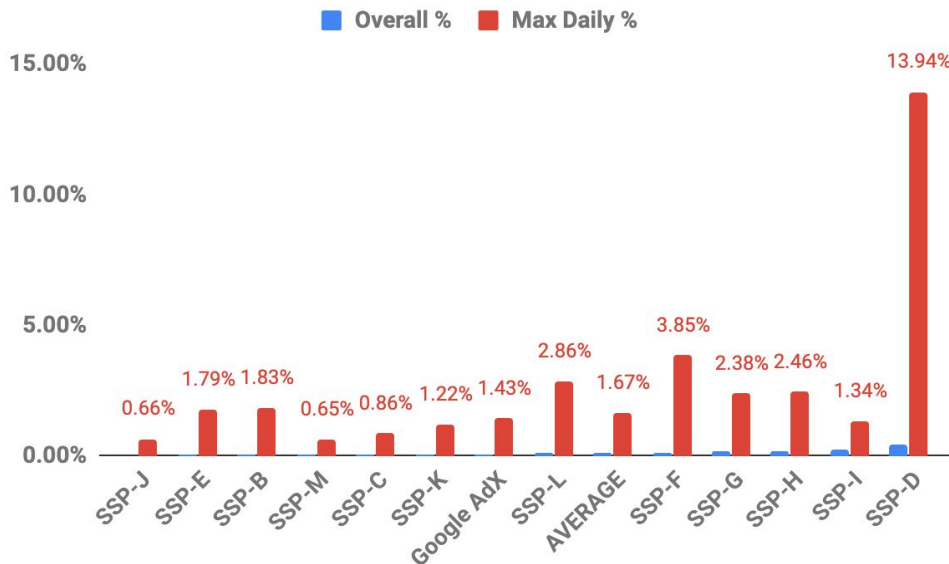
SSP-I and SSP-D were the worst performing SSPs in Q3, but both saw significant improvement vs. Q2. Despite this, SSP-D is over 20x as likely to deliver a Malicious ad than the best performer.



Malicious ads remain a fairly concentrated problem: Nearly 60% of impressions came from just 3 of the 75+ SSPs monitored by Confiant. Most alarmingly, a single SSP was responsible for 30% of malicious ad impressions.


Malicious ads remain a fairly concentrated problem: Nearly 60% of impressions came from just 3 of the 75+ SSPs monitored by Confiant. Most alarmingly, a single SSP was responsible for 30% of malicious ad impressions.

Of course, 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.



Malicious ads remain a fairly concentrated problem: Nearly 60% of impressions came from just 3 of the 75+ SSPs monitored by Confiant. Most alarmingly, a single SSP was responsible for 30% of malicious ad impressions.

Our data shows that SSPs are likely to be hit with at least one major attack every quarter. When these attacks occur, how bad can we expect things to really get? It turns out quite bad, as Chart 4 illustrates. **Top performers SSP-E and SSP-B saw peak levels come in at 45x their overall average.** The worst performer on both peak and average rate, SSP-D, had a peak daily Malicious ad rate of 13.94%, showing just how high these rates can get in the midst of an attack.

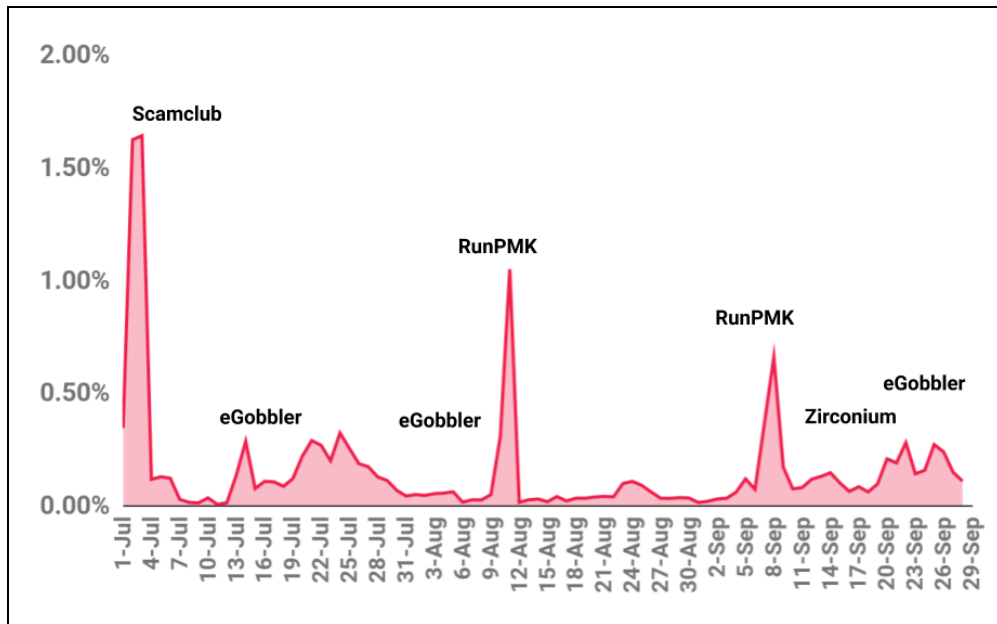


*faMk PC fZMIZ?Zyuni yZ'Ab Z'jb ž MyAn\_{nu'@y{n'2 Nj, MZ° {(Mly*

Similarly, it's important to understand how quickly an SSP responds to Malicious ads when an attack is underway. On this measure, we see huge disparities between the best performers and the worst. Not surprisingly, the two SSPs that had the highest rate of Malicious impressions, SSP-D and SSP-I, were among the slowest to respond to attacks.

## Major Threat Groups Active in Q3

Most attacks stem from a fairly small number of highly sophisticated threat groups that specialize in exploiting the fragmented adtech ecosystem. We have [written extensively](#) on many of these actors, even going so far as to assign monikers to them to facilitate tracking over time. In Q3, four of these threat actors were responsible for most major attacks. Below we describe each of those threat actors and the characteristics of their Q3 attacks:



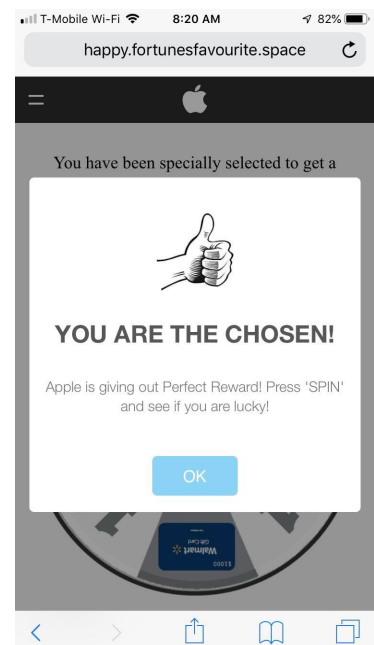
faM(B@ {{(Ml M(b)† MIXk Mnx{axZMlMk(nxyrfZx> UÜÜÜä

### GMa W Vs

**Date(s) of peak activity:** July 1, 2019

**Notable characteristics of attack:** Consisted of a highly targeted attack on a single DSP.

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 inevitably spill beyond any anti-malvertising gatekeeping.



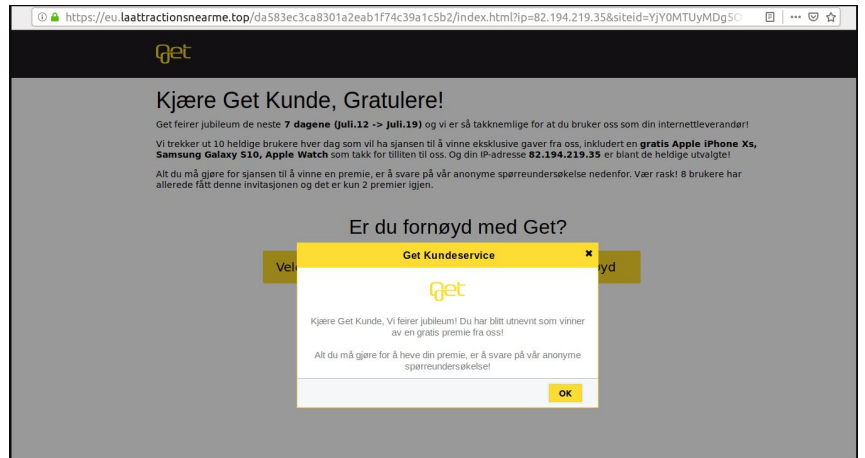
## Y; cWYf

### Date(s) of peak activity:

- July 16, 2019
- August 7, 2019
- September 28, 2019

**Notable characteristics of attack:** This 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.

Their Q3 attacks targeted desktop computers, mainly running Windows, with high concentrations of users in Italy, Spain, and Scandinavia. Our researchers found that even when publishers set up iframe sandbox permissions optimally, a pop-up could be spawned when the user tapped on the parent page. Confiant reported this vulnerability to the Webkit team on August 7, and it was fixed in iOS 13. Over the course of its various iterations on Chrome and Webkit (Safari), the eGobbler attack [infected over 1 billion ads](#).

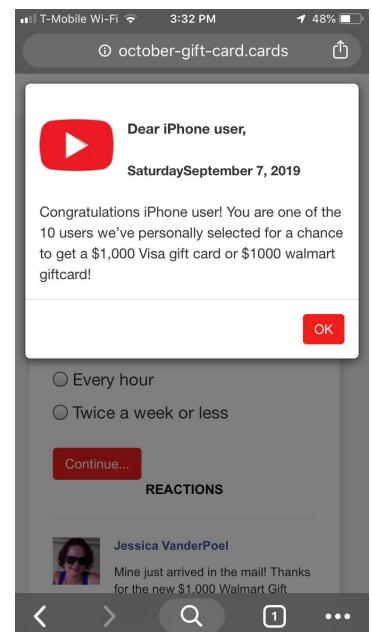


## Fi bDA?s

### Date(s) of peak activity:

- July 12, 2019
- August 11, 2019
- September 9, 2019

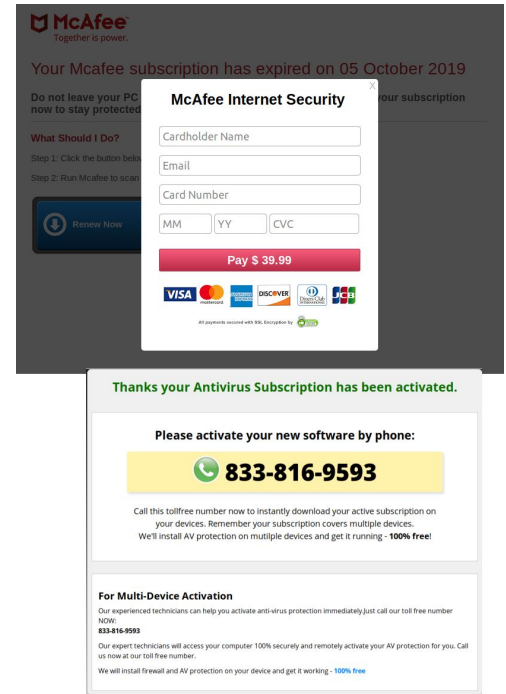
**Notable characteristics of attack:** Notably, all three redirect attacks came through Google DV360, an unusual vector due to Google's considerable malware defenses, and then spread rapidly across multiple SSPs as well as AdSense. Focusing on mobile traffic (both iOS and Android), these combined attacks ran globally across 212 countries. At its peak, RunPMK controlled up to 2% of the overall display demand and its impact was seen on virtually every RTB-monetized site.



## Malicious Ads

**Date(s) of peak activity:** mid-September.

**Notable characteristics of attack:** 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.



## Why has the Rate of Malicious Impressions Declined?

An astute reader will note that the rate of Malicious Ad impressions has declined markedly over the course of 2019. How should we interpret this trend? Has the industry really improved that significantly over the last nine months? The answer is a complex one that includes factors related to both the methodology Confiant uses to compile this report as well as industry-wide trends. We lay out our thoughts on these factors below:

**Attacks vs. impressions:** This report is designed to measure the frequency with which bad impressions are delivered to a user's browser and blocked by Confiant's client-side verification solution. We believe the most meaningful way to represent this data is by showing what percentage of impressions would have been impacted by bad ads had our solution not been in place, as this most closely tracks the impact to users. However, it does not track the number of unique threats or incidents present in the ecosystem over time. An individual threat can lead to 10 impressions or 1 million impressions; obviously, it's the latter one we should care about. If, by working with SSPs and DSPs, Confiant can interrupt an attack in progress and prevent the ads from making it to the user's browser, it will appear as if the number of bad impressions has been reduced. However, the underlying threat to the ecosystem remains the same — it was merely resolved upstream. The distinction between attacks and impressions shows up in our data: **the volume of individual attacks increased in August and September despite the decline in bad impressions.** There are many opportunities to divert a bad ad before it's delivered to a user, and we must make sure we as an industry are utilizing all of them.

**Selection bias:** To assemble this report, Confiant uses a sample of over 120 billion impressions collected across our publisher partners. While this is a robust sample in terms of size and diversity, it's inherently biased by its inclusion of Confiant customers only. In an ideal study, Confiant would pull client-side data from a completely random selection of websites and users. Unfortunately, this is infeasible in practice as our code would need to appear on either every site on the Internet or in every user's browser. The lack of randomness introduces selection bias:



Publishers who have taken the step of implementing a creative verification solution like Confiant's can reasonably be expected to be more quality-conscious than the average website. And in fact, our publishers frequently make use of the extensive data we provide to cull poorly performing SSPs and institute other protective measures such as blocking specific DSPs and seat IDs. The end result is that the volume of bad ads that make it to the user's browser declines as publishers put these protections in place. Given these factors, this report will tend to *understate* the frequency of bad ads compared to a general survey of the internet. Still, we believe the Demand Quality Report has value both as a measure of the progress that's possible when collaboration occurs AND as a way to highlight the ongoing disparities in performance across different SSPs.

**Industry collaboration:** We strongly believe that Malicious ads are an industry-wide problem that warrants an industry-wide solution. The burden of stopping bad ads shouldn't fall on publishers alone. To that end, Confiant began offering a free data solution to top SSPs in Q2 to assist them in identifying and shutting down the largest malware attacks. SSPs have enthusiastically adopted this tool, as it allows them to shut down major attacks quickly and protect their publishers from the fallout. In addition to this work with SSPs, Confiant has collaborated with major browsers to eliminate vulnerabilities that allowed malvertisers to bypass normal sandboxing restriction to launch pop-ups and redirects. Over the course of 2019, Confiant research has led to the elimination of major vulnerabilities in Webkit browsers ([CVE-2019-8771](#)), Chrome ([CVE-2019-5840](#)), and Opera.

**Seasonal trends:** In 2018, we saw a reduction in Malicious ads in Q3, only for the rate to explode in Q4. We anticipate a similar trend in 2019, as malvertisers seek to exploit reduced staffing levels at publishers and platforms over the holidays.

A host of industry-wide factors also continue to change the game for malvertisers:

The near-universal adoption of ads.txt on top sites, which is driving arbitrage and unauthorized resale out of the market and foreclosing IBV opportunities.

Increased vigilance on the part of the SSPs when it comes to ad quality issues given the highly competitive, and increasingly commoditized, nature of the space.


Industry initiatives like TAG's Certified Against Malware program, which have increased general awareness around the threat of malware and galvanized efforts to combat it.

Better coordination between publishers and platforms, which narrows the brief window of time that malvertisers have before their exploits are detected and removed.

We celebrate the decline in the rate of Malicious ads on our partners. It's the hard-fought result of the efforts of the entire ecosystem. But the problem of Malicious ads is by no means solved, and any lessening of pressure on malvertisers will allow the threat to come roaring back.

## Malicious Rates Increase Markedly on Weekends and Holidays

It will come as no surprise to veterans of programmatic advertising that Malicious activity increases on weekends and holidays:



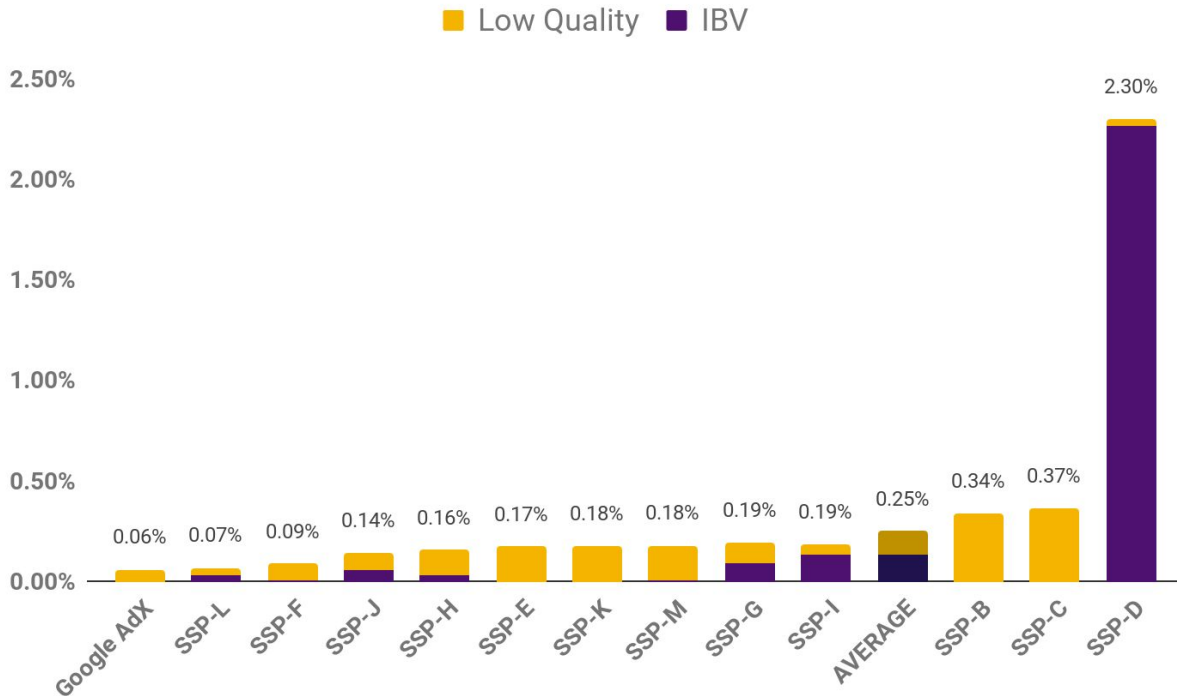
*Malvertising is a significant risk to advertisers and consumers. Malicious activity is more likely to occur on weekends and holidays.*

Malvertisers target these times to take advantage of reduced staffing levels and slower response. What may come as a surprise is the extent to which the risk of malicious activity varies by day of week and on special days like holidays. Looking at a whole year's worth of data, we found that **an impression served on Sunday is 54% more likely to deliver a Malicious payload than one served on an average day.** Even more alarming, **an impression served on a holiday is 141% more likely to be malicious.**

## Low Quality Ads

The final measurement category, Low Quality ads, is derived from a diverse set of rules that publishers can elect to activate on the Confiant platform. These rules correspond to ad behaviors that have one feature in common: they disrupt or impair the user experience. Examples include autoplay audio, autoplay video, pop-ups, and In-Banner Video.

In the past, we have reported separately on In-Banner Video and Low Quality ads. Given that both sets of issues lead to a disruptive and frustrating experience for users, we will be combining these categories in this and future reports. The chart below shows how this combined measurement varies across the top 13 SSPs. The overall frequency of Low Quality ads declined slightly from Q2 to Q3, falling from 0.25% to 0.22%. The standout performer — for all the wrong reasons — was SSP-D, which saw its rate climb from 0.29% in Q2 to 2.30% in Q3, driven by an explosion in In-Banner Video. Conversely, SSP-M improved their performance from 0.93% in Q2 to 0.18% in Q3.



With the exception of SSP-D, In-Banner Video has largely been driven out of the ecosystem. To explain why, it's important to understand that almost all IBV originates from an arbitrage opportunity. CPMs differ enormously between Video and Display. This creates an arbitrage opportunity wherein an unscrupulous advertiser can buy a display impression, insert a video player, and then resell the impression as a video opportunity at a much higher price to an often unsuspecting buyer. In the past, some SSPs allowed resellers to exploit this arbitrage opportunity to the detriment of their buyers, who

With the exception of SSP-D, In-Banner Video has largely been driven out of the ecosystem. To explain why, it's important to understand that almost all IBV originates from an arbitrage opportunity. CPMs differ enormously between Video and Display. This creates an arbitrage opportunity wherein an unscrupulous advertiser can buy a display impression, insert a video player, and then resell the impression as a video opportunity at a much higher price to an often unsuspecting buyer. In the past, some SSPs allowed resellers to exploit this arbitrage opportunity to the detriment of their buyers, who


think they are buying a traditional video impression, and their publishers, whose legitimate video inventory is being devalued by this practice.

Industry initiatives such as ads.txt have closed this arbitrage opportunity by making it plain to all market participants which providers are authorized to sell a publisher's inventory. Combined with related moves by the DSPs to optimize their supply chains, increased adoption of ads.txt is driving out unauthorized resellers from the ecosystem and taking IBV with it.

The incidence of IBV ads is highly concentrated, with over 50% coming from just 3 providers.

## Q3 Rates by Country

Our data shows significant variation in the rates of Malicious, In-Banner Video, and Low Quality ads by country:

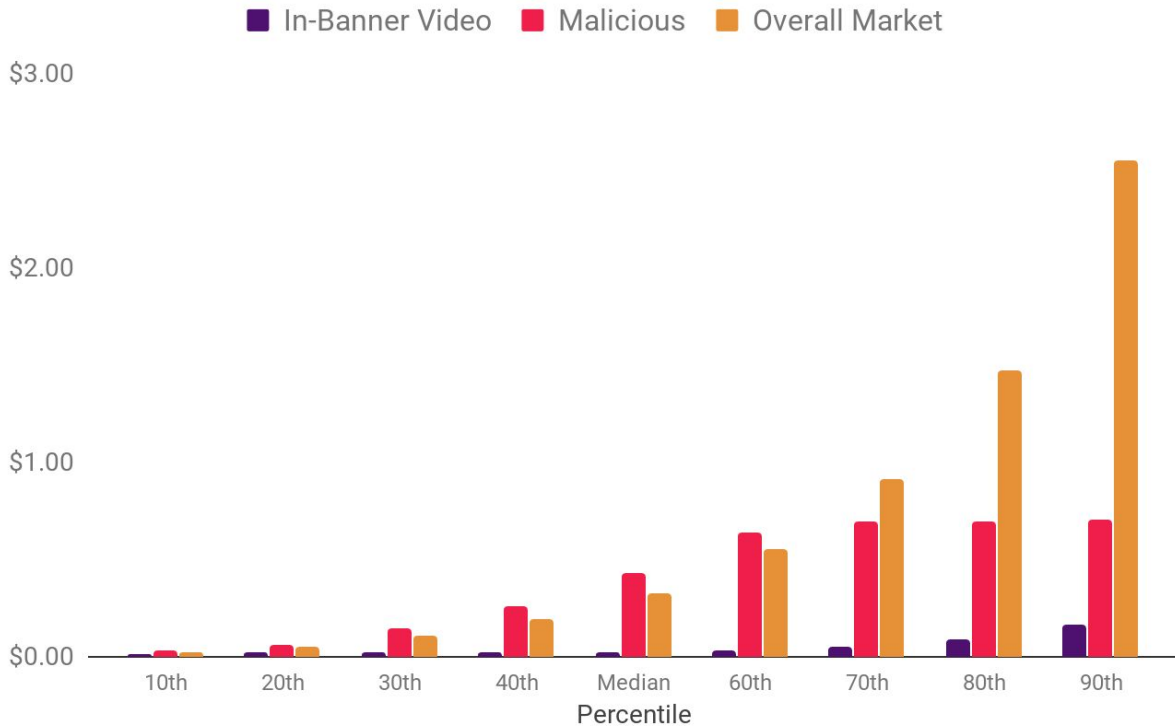


.....t aM'â@3nx(a" k ZxMI MIX"/ xnuZMI Mx w/ Nq+ xMZYb >Ü

As in past quarters, European markets saw far higher rates of Malicious ads than the U.S., but a lower rate on other issues. In Q3, the rate of Malicious ads increased significantly in Italy, France, and Spain, while declining modestly in Germany and Great Britain. The variety of rates by country exemplifies how malvertisers continually shift their campaigns and targets to remain under the radar.

## Q3 US CPMs

Below is the CPM distribution in Q3 for Malicious ads, In-Banner Video ads, and the market as a whole:



### Key Takeaways

While the data clearly shows that a strategy of raising floors can be effective at blocking IBV, that strategy becomes ruinous when applied to Malicious ads. That’s because CPMs for Malicious ads match — and in some cases exceed — those of the overall market until we reach the 80th percentile. To provide a specific example, setting a floor of \$0.70 would block 90% of Malicious ad impressions and nearly 100% of IBV ad impressions; however, it would also block 65% of clean ads. The resilience of Malicious ad CPMs demonstrates just how lucrative the act of malvertising can be: the malvertisers are quite willing to spend in line with general advertisers to obtain the audience they value.

## Conclusion

The results of the Q3 Demand Quality Report demonstrate the continued relevance of the problem of bad ads while also giving hope for the future. The frequency of bad ads remains unacceptably high: we found that nearly 1 in every 250 impressions is dangerous or disruptive to the end user, which equates

Nearly 1 in every  
250 impressions is  
dangerous or  
disruptive to the  
end user

to 4 billion malicious or disruptive impressions a month across the entire industry. The rapidly evolving tactics of top threat actors, who skillfully identify and exploit browser vulnerabilities to generate redirects even in the presence of iframe sandboxing, show that this is not a problem that is going away any time soon. However, we are encouraged by the continued decline in the rate of bad ads on Confiant publishers, which demonstrates that there are effective mitigation methods, both in terms of technology and partner selection, available to those who wish to use them.

As in past reports, the results reveal a vast disparity between the best and worst performing SSPs when it comes to ad quality. Among top SSPs, the worst performer is almost 20x as likely to deliver a bad ad compared to the best. Likewise, the data shows that many quality issues are highly concentrated:

- Over 50% of Malicious impressions came from just 3 providers.

- Over 50% of Low Quality ad impressions came from just 3 providers.

- A single top SSP had the dubious honor of appearing among the worst 3 providers for both Malicious and Low Quality ads. This SSP is in the top five for impression volume.

But even strong performers can have off quarters: Google fell from the top spot in preventing Malicious ads in Q2 to seventh place in Q3 (though still remains an above-average performer).

Publishers need to understand the risks of bad ads as well as the techniques to mitigate these risks. Our Demand Quality Report will continue to highlight both as we fight as an industry to reduce the scourge of bad ads.

## About Confiant

Confiant is a cyber security company that came out of a recognition that the world's most sophisticated advertisers aren't Verizon or P&G, but criminals using the industry for their own, selfish ends. These criminals are hijacking programmatic advertising and giving publishers a bad name.

Confiant protects the reputation, revenues, and resources of publishers and platforms with always-on anti-malware software that verifies desktop, mobile, and video ads. Our sole focus is on helping advertising platforms and publishers rid the world of malware. This focus enables us to evolve quickly and meet our clients' needs for defeating the bad actors trying to undermine the industry.

We were the first to come to market with a technology that does not just detect malicious activity, but actively blocks it. We believe in the intelligent application of technology to fight back and make digital media safe for everyone.