

A man with glasses is shown in profile, looking at a computer screen. The screen displays lines of code, likely JSON, which is the subject of the whitepaper. The background is dark and blurred, focusing attention on the man and the code.

AN INDUSTRY WHITEPAPER

The Urgent Need For Buyers.json

Shutting the door on malvertising

Introduction	3
What is buyers.json?	4
Composition of the Buyers.json file	5
How will this work in practice?	8
Benefits to the buy side	9
What Buyers.json enables	11
Why support from premium publishers is so critical	12
How you can contribute	13

Introduction

Digital advertising has long been plagued by twin threats: ad fraud and malvertising. The industry has made tremendous strides on ad fraud with initiatives like ads.txt, sellers.json, and the SupplyChain Object. However, it still struggles to deal adequately with malvertising. Indeed, malvertising and other ad quality issues have long been overshadowed by fraud largely because the victims are so diffuse: the everyday users of websites. Users do not have a seat at the table in industry discussions, and they do not control large budgets that can be used as leverage to force change.

The fact is, we as an industry have not been good stewards of the ad experience. By allowing malvertising and other disruptive ad experiences to propagate unchecked, we have driven a large subset of the public into taking extreme measures such as ad blockers. If we are to make the argument that publishers and users are engaged in a fair tradeoff of value, we must be able to state in good faith that we are taking appropriate steps to reduce risks for users. Yet the industry continues to fall short here. Today, the door to the ad tech ecosystem is largely open: with a creative and a credit card, anyone can buy an ad and often go live with minimal vetting and oversight. Advertisers are attracted to digital advertising for the enormous reach it offers, its advanced targeting capabilities, and the ability to measure results quickly and iterate. These same factors make digital advertising a powerful attack vector for malvertisers.

One of the chief challenges with eradicating malvertising is buyer identification. Malvertisers take advantage of the highly fragmented nature of the ecosystem by exploiting the weakest links and jumping frequently from DSP to DSP. As soon as they are found out on one DSP, they seamlessly shift their campaigns to a new one. Publishers and SSPs seldom have insight into the identities of buyers and therefore cannot relay the information to the DSP that would allow them to block a known bad entity across all the access points at once. The end effect is whack-a-mole, with the same entity appearing again and again across different DSPs.

The answer is greater buy-side transparency, with one solution rising above all others in terms of efficacy and ease of implementation: buyers.json.

What is buyers.json?

Buyers.json is a simple mechanism to allow DSPs to publicly share the names and identifiers of the buyers they represent, facilitating quick identification of threat actors when attacks occur. It acts as the mirror image of Sellers.json, which allows SSPs to share the identities of the publishers they represent. Buyers.json would bring transparency to the buy side in the same way that prior initiatives brought transparency to the sell side. This is important not only as a matter of fairness (parity between buy-side and sell-side) but more critically, as a necessary step to protect users from malvertising.

Buyers.json follows in the path set by the industry's very successful sell-side transparency efforts. As Sellers.json boosted transparency into the identities of sell-side entities, Buyers.json does the same for the buy side. Disclosure of buyer names solves two major issues for the industry:

- It supports identification of malvertisers and other threat actors across DSPs
- It facilitates blocking of identified threat actors across DSPs

A Buyers.json spec is currently being drafted by an IAB Tech Lab working group composed of a diverse array of marketplace participants.

Unlike other potential solutions considered by the IAB Tech Lab, Buyers.json offers open access with a simple, human-readable format — literally anyone with a web browser can view and interpret the file. No setup of complex APIs is required, nor is it necessary for a publisher to have a relationship with their buy-side counterparts. Simply put:

*BUYERS.JSON DEMOCRATIZES
ACCESS TO INFORMATION THAT
IS CRITICAL TO THE SAFE AND
EFFICIENT FUNCTIONING OF
DIGITAL ADVERTISING.*

Composition of the Buyers.json file

The specification for Buyers.json is still being finalized by IAB Tech Lab. But in all likelihood, it will largely mirror the existing Sellers.json spec, but with fields for buyer ID, buyer name, and buyer type replacing their seller ID, seller name, and seller type, respectively. The buyer ID would match the ID that appears in the bidresponse.seatID property in OpenRTB and the proposed DemandChain object, allowing these values to easily be tied together. The buyer type attribute would indicate whether the buyer is acting as an advertising or an intermediary.

DSPs will post the buyers.json file on their root domain. Because Buyers.json is in a simple, human-readable format, hurried adops and compliance teams trying to pinpoint the source of an attack will be able to benefit from the information in these files without the support of APIs, specialized tools, or a large technical staff. More sophisticated parties such as SSPs will be able to crawl these files on their own schedule and eliminate manual processes for exchanging buyer information with DSPs.

Why now?

Our research indicates that at least 1 in every 150 impressions¹ is malicious or disruptive to the user experience. When taken in the context of the entire industry, the resulting numbers are truly staggering. We estimate that 20% of user sessions are impacted, meaning billions of attempts to compromise users take place each month. Faced with these numbers, it's no surprise that users flock to ad blockers.

The industry has erected defenses against this constant assault, including certification processes (TAG Anti-Malware), creative verification solutions (Confiat and others), threat-sharing exchanges, and scanning best practices. However, the highly fragmented nature of digital advertising presents the perfect target for

¹ Q4 2019 Confiat Demand Quality Report

malvertisers. To enter the ecosystem, bad actors need only find the weakest link in the chain. Indeed, malvertisers are masters of business development -- they are shrewd at convincing DSPs of their bona fides and their willingness to spend right now (the proverbial 5pm Friday high-dollar malvertising campaign has an element of truth). They exploit limited staffing on weekends. And most of all, they exploit identity. Publishers are often not aware of who is on the other side of a programmatic ad buy. Is it a direct advertiser? A massive agency holding company? Or a fly-by-night "agency" whose sole intent is to find and exploit gaps in security? And even if the publisher is able to identify the source of the attack, how can they ensure that entity is not present on other DSPs that have access to their inventory? The core of Buyers.json is answering this question.

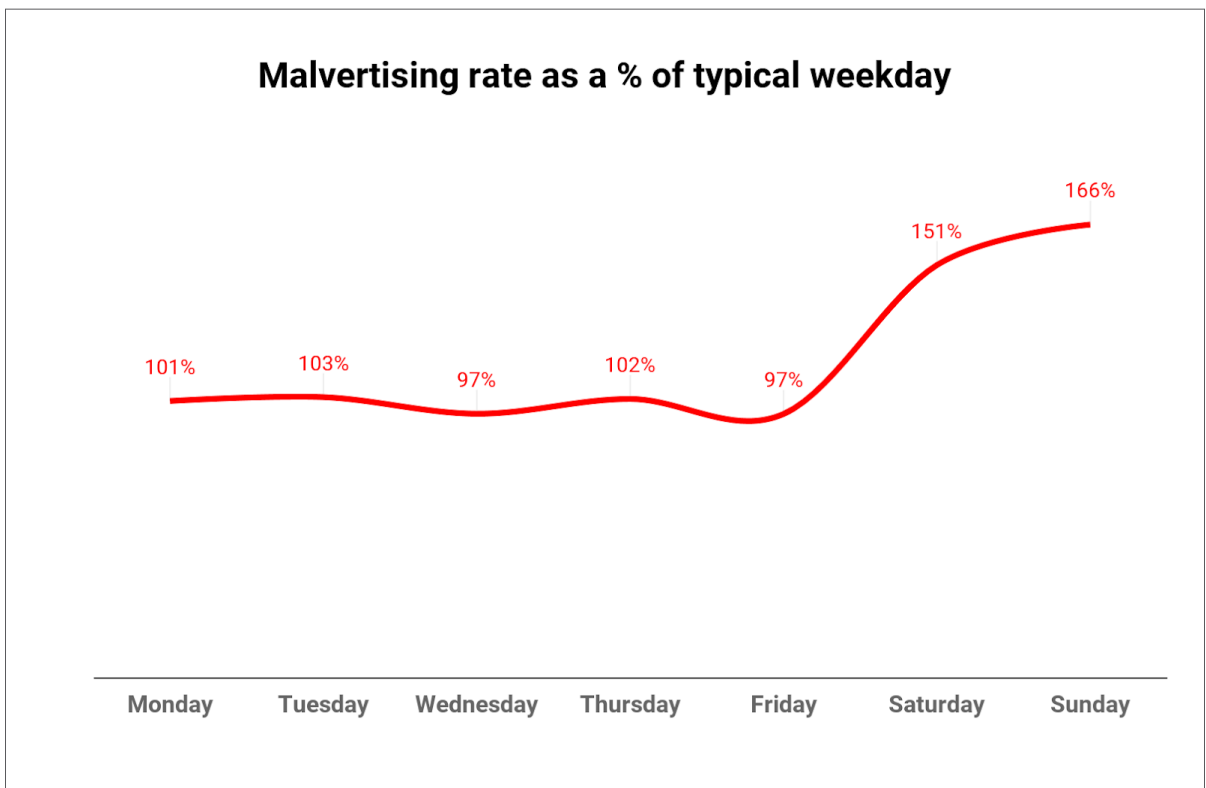


Figure 1. It's not just in your head: Malvertisers really are working for the weekend.

A dirty secret of digital advertising is that when a publisher runs an ad, they are allowing a third party (and often an unknown one at that) to not only inject content onto the page but to run code. Now most of this code is entirely harmless (e.g. viewability measurement, brand safety checks, analytics). However, a small percentage is pernicious and actively seeks to compromise the user through malware, forced redirects, phishing attempts, and criminal scams. When publishers

identify the entity behind an attack, their natural reaction is to attempt to block it. So they might ask their SSP to block the buyer from accessing their inventory in the future. SSP dutifully blocks the buyer ID behind the attack, which corresponds to one account on one DSP. Whether that buyer has a presence on other DSPs is almost always left unaddressed for the simple reason that in many cases, the SSP also does not know who the entity is. That same entity could have multiple accounts on a DSP or even scores of accounts across multiple DSPs. Because the supply side lacks transparency into buyer identity, it is nearly impossible today for a publisher to block a buyer in all its forms (both within and across DSPs) from their inventory.

*PUBLISHERS SHOULD KNOW
WHO HAS ACCESS TO THEIR
INVENTORY AND THEIR USERS.*

Another reason why this is so important now is fairness. Over the past three years, buyers have gained tremendous visibility into the inner workings of the sell-side through ads.txt, Sellers.json, and the SupplyChain object. While much of this data is used for the commendable purpose of fraud prevention, it's also been used extensively by advertisers and DSPs for supply-path optimization (SPO) and other purposes. This tremendous increase in transparency for the buy-side has not been matched by any such gains for the sell side. This has worsened the already asymmetric relationship between buyers and sellers. Buyers now have the ability to determine how best to access your inventory; publishers do not possess the corresponding data to determine how best to sell their inventory. This skews what should be a neutral transaction in favor of buyers. Buyers.json is the first step toward restoring balance between buyers and sellers, which is good for the industry and just the right thing to do.

Finally, transparency itself is a virtue. Digital advertising—and in particular programmatic advertising—has long been criticized for being opaque. This opacity has allowed bad behavior to flourish, whether in the form of malvertising, brand safety issues, privacy violations, or disruptive ad experiences. Extending the transparency that has been achieved on the sell-side to the buy-side will for the first time allow a bad ad to be tracked from its source to its destination, and through all points in between. This level of transparency will flush out bad actors who previously could hide behind a wall of obfuscation.

How will this work in practice?

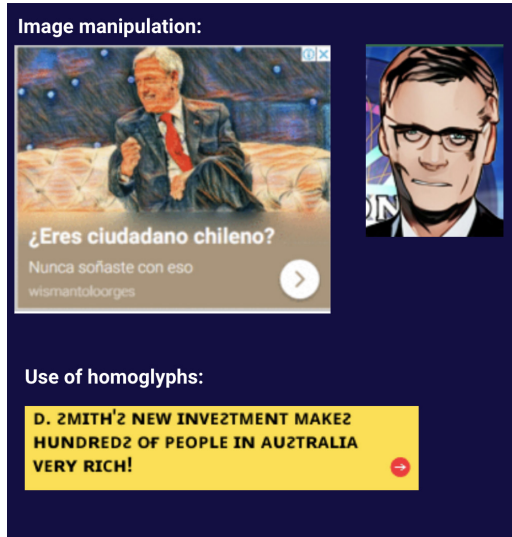


Figure 2. Some tricks of the trade to evade detection. Images are manipulated and homoglyphs might be incorporated into text so it is harder for facial recognition and OCR to identify.

A reasonable question is whether Buyers.json will really help in identifying all of a malvertiser's accounts across DSPs. Won't malvertisers just set up scores of accounts all under different names? Sure, they'll try, but it's not quite that easy. To understand why, it's important to understand who malvertisers are and how they infiltrate the ecosystem. The truth is that most malvertising comes from a fairly small number of global criminal organizations. These organizations are not three hackers in a basement in Cyprus. They are large criminal enterprises who become highly specialized in different methods of attack. Some exploit vulnerabilities in browsers (Confiant research has led to the elimination of major vulnerabilities in Webkit browsers (CVE-2019-8771)

and Chrome (CVE-2019-5840)). Others employ image manipulation and use homoglyphs to trip up machine classification of creatives. Almost all attempt to evade detection by fingerprinting malware scanners and deploying their payloads only to specific users and regions. In short, they are extremely innovative and highly incentivized to overcome any challenges that the industry puts in their path.

But in order to work with top-tier DSPs and gain wide access to the ecosystem, malvertisers need to look like real businesses. This requires articles of incorporation, bank account numbers, websites, email addresses, fake LinkedIn profiles, and all the rest.

The creation of these assets takes time, money, and effort, and no malvertiser has an infinite supply of fake personas they can tap into. They must instead invest significantly in each profile and use it as long as it remains possible to do so. The

fact that it's nearly impossible today for publishers and SSPs to identify an entity across all DSPs makes it easy for malvertisers to keep this game up for weeks or months. Malvertisers can burn a profile on one DSP, but easily stay live on another. Buyers.json is designed to solve this problem by providing a simple mechanism to identify a bad actor across multiple DSPs, allowing the industry to respond far more quickly and completely to attacks. Beyond reducing time to resolution, Buyers.json will dramatically increase the costs for malvertisers and dissuade additional attacks. Lastly, Buyers.json will also facilitate threat-sharing by providing a common identity to which malvertisers can be tied.

Benefits to the buy side

We've explained why Buyers.json is vitally important to the sell side. But how will it benefit buyers and DSPs, who will ultimately bear the cost of implementing the standard? While the sell side clearly loses the most to malvertising, legitimate advertisers are adversely affected as well. Malvertisers frequently steal brand assets (e.g. logos, design elements) and creatives of major companies to give their campaigns an air of legitimacy. A user affected by a phishing campaign that uses a legitimate advertiser's branding might associate that bad experience with the brand owner, not the criminal enterprise that has hijacked the brand's assets. Such misrepresentation could also lead to legal action against the buy side. Sophisticated malvertisers have lately adopted nationally recognized celebrities as brand ambassadors to sell their wares. The scammy campaigns that misappropriate a celebrity's name or image could also place the DSP who delivered it at significant legal risk.

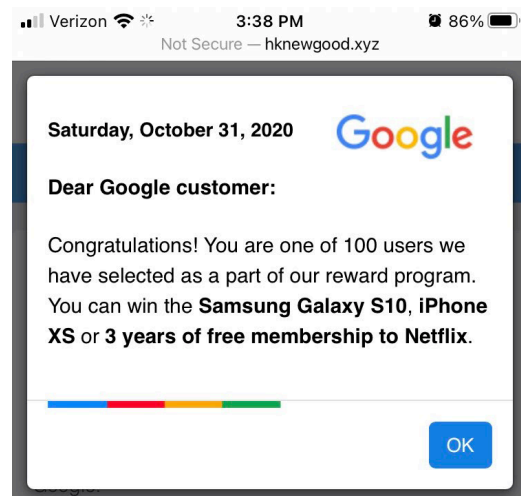


Figure 3. Malvertisers commonly misappropriate brand elements to give their campaigns an air of legitimacy.

Malvertisers also compete with legitimate advertisers for publisher inventory, and their actions poison customer perceptions about digital advertising. Much like the well-known “tax” from ad fraud, the false competition from malvertising imposes an invisible levy on legitimate advertisers by driving up the cost of impressions. Worse, the negative experiences created by malvertising drive user adoption of ad blockers and thereby reduces reach. As an industry, we’ve already lost access to more than 25% of users who have adopted ad blockers; it’s imperative that we stave off further losses by protecting vulnerable users from dangerous attacks.

Finally, our failure to address a persistent threat to the public invites regulatory action that is likely to fall primarily on buy-side platforms. As an industry, we have a limited window to implement a solution to malvertising before one is imposed on us by regulators who have limited knowledge of digital advertising. Considering the regulatory burden of GDPR and CCPA, we should not repeat the mistakes of the past and squander this opportunity to chart our own course.

In total, the costs imposed by malvertising on brands and DSPs in the form of increased CPMs, damage to brand reputation, loss of reach, and regulatory risk far outweigh the cost of adopting Buyers.json. While the sell side might benefit most from Buyers.json, the buy side also stands to gain from its implementation.

What Buyers.json enables

Buyers.json is the first step in boosting buy-side transparency, just as ads.txt was a first step in boosting sell-side transparency. But it will also be the foundation upon which future initiatives will be built. By solving the problem of cross-platform buyer identity, Buyers.json will allow the industry to turn its attention to new efforts like:

- DemandChain, a new object within OpenRTB that would allow sellers to see all parties that were involved in buying the creative. Just as buyers.json serves as the counterpart to sellers.json, DemandChain is envisioned as the buy-side complement to the SupplyChain.
- Standardizing methods for disclosing buyer information on the client-side and through header-bidding frameworks like Prebid. In an ideal world, every creative would have a “calling card” that clearly identifies who bought inventory and who delivered the ad to the page. With this information in hand, Buyers.json would be the directory to which a publisher would turn to look up the buyer and connect it to specific accounts across DSPs.
- Creating a centralized registry of buyers and exchange mechanism for reputation signals.
- Simple, standardized ways to manage and communicate ad quality preferences across SSPs.
- An industry certification program for ad servers.

Why support from premium publishers is so critical

Collective action is hard. Even when the benefits are clear to all parties, coordination costs and the presence of slightly misaligned incentives can slow adoption of new standards. Successful efforts can take years. In fact, the industry's sell-side transparency initiatives, which are widely regarded as a major success, took years of discussion, planning, and coordination to come to fruition. Ads.txt was introduced in 2017, followed by app-ads.txt the following year. Sellers.json and the SupplyChain Object arrived only in 2019. In each of these cases, major advertisers or DSPs provided support or applied pressure at critical junctures. Ads.txt likely would not have been successful without the vigorous support of Google and major advertisers like P&G. Similarly, Sellers.json and SupplyChain Object would never have been adopted so quickly without the strong advocacy of The Trade Desk. But in these cases, advocates of transparency had a major cudgel: as advertisers, agencies, and DSPs, they controlled the flow of dollars into the industry. If they put their foot down, most SSPs and publishers had no choice but to fall in line.

Can we expect similar results for buy-side transparency efforts without that sort of leverage? Probably not. While DSPs have nothing against efforts to combat malvertising (in fact, many DSPs are quite engaged and supportive), they are unlikely to reveal their client lists or schedule dev work without strong pressure from the sell side. The problem is that no single entity has the clout to force adoption of Buyers.json. It's therefore imperative that the sell side lock arms and demand that DSPs comply for protection of users and the good of the industry. We want every Comscore 250 to commit to this fight. Publishers need to realize their power: Advertisers need you and need your users. As an individual publisher, advertisers and DSPs might be able to ignore you. However, as a group you set the terms by which the buy side accesses your inventory.

How you can contribute?

In addition to this white paper, Confiant is building a microsite to spread the word about Buyers.json. We believe it will be a game-changer for buy-side transparency, accelerating the industry's efforts to eradicate malvertising and restoring some balance to the relationship between buyers and sellers. But acting alone none of us can convince the buy-side to support Buyers.json. Instead, we need to create a movement:

We need premium publishers and SSPs to come out in support of Buyers.json and to publicly commit to favoring sources of demand that participate.

We will highlight these commitments on our microsite and work to educate our partners and the wider industry.

While this will initially focus on the supply side, we believe we have a message that will resonate with advertisers as well.

For more information visit the microsite at:

buyersdotjson.com