ExchangeWire deep dive

Out-Evolving Ad Fraud: How Ad Tech Can Tackle Current and Emerging Threats

In association with

White Ops

 www.exchangewire.com research@exchangewire.com

CONTENTS

INTRODUCTION

Over the last couple of decades, digital marketing has become the lifeblood of virtually every marketing organisation in the world. But the speed of innovation and development of new marketing techniques has left openings for fraudsters to swoop in to try and capture a piece of the enormous spend.

The unfortunate truth is that many aspects of digital marketing have been pervaded with fraud. With the internet came new opportunities for advertising, and with new opportunities for advertising came new opportunities for fraudsters to game the system. With the speed and volume of advertising transactions at levels that were unprecedented just 20 years ago, malicious actors have enormous opportunity to exploit systems' vulnerabilities.These actors steal and waste marketers' ad spend on lead generation, digital and programmatic advertising, and organic social and search campaigns, impacting the organisation's tech stack and their budgets.

At a time when marketers are under increasing pressure to focus on metrics and conversions, as e-commerce pushes the economy towards digital, bots are entering CRM systems and data management platforms, skewing results, and disrupting the online advertising ecosystem. Not only are these bad actors stealing marketers' money, but they are also damaging their predictive models, paving the way for future ad spend to be wasted.

Despite knowing that ad fraud is very real, for many in the industry, how it is carried out remains an abstract concept. There are multiple techniques that bad actors implement to siphon off valuable budgets, and being able to recognise and understand these techniques is imperative to shutting them down. We have outlined some common methods employed by fraudsters to run their operations.

White Ops has produced this Deep Dive special, in partnership with ExchangeWire, analysing the current landscape of ad fraud and highlighting the evolving tactics of bad actors. White Ops is the global leader in collective protection against sophisticated bot attacks and fraud. White Ops verifies the humanity of over 10 trillion online interactions every week to protect enterprises from across the globe from some of the most sophisticated bot attacks.

In this DeepDive special:

01. An overview of common techniques used to commit ad fraud

02. What the industry can do to tackle fraudsters

03. Emerging threats in new channels

04. Industry players' perspectives into how to address the risk of ad fraud

PART 1: COMMON AD FRAUD TECHNIQUES

App spoofing

App spoofing is the practice of an app sending false bundle ID information to the ad exchanges, pretending to be something it's not. It's a fairly common, but also fairly manageable, way for fraudsters to make money. App-ads.txt is an initiative from the IAB Tech Lab intended to prevent misrepresentation by requiring app developers to include a list of which monetisation partners are permitted to sell inventory within the app.

Out-of-context ads and navigation

Fraudsters often utilise out of context adverts and forced pop-up units to maximise false impressions, usually to the great frustration of end-users. Out-of-context placements typically involve rendering ads, including native and interstitial, from major ad-networks, while out-of-context navigation launches false intents to URLs received from the command and control server (C2).

A fraud operation that was carrying out the out-of-context ads and navigation techniques was uncovered by White Ops recently. Detected by the Satori Threat Intelligence and Research Team, the fraudsters were running irrelevant ads across native and interstitial, launching out-of-context navigation intents, and removing the icons of the 38 beauty apps it was operating across to make it harder for users to uninstall them.

Whilst this operation was a classic case of the out-of-context fraud technique, it also exemplified how quickly fraudsters can adapt their strategies in order to continue after being discovered. This included some unusual behaviour by those behind the scheme that could potentially point to a lesser-known, or possibly even new, ad fraud technique.

With data arriving in inconsistent forms from disparate sources, it's not shocking that it takes so much time to organise and interpret. But, as anyone can see, this process is simply inefficient — not only is manual data integration incredibly taxing, but putting the onus of collating copious amounts of data onto employees opens up the risk for all too natural human error. Furthermore, the amount of time it takes to complete manual integration means that even when these different data sets have been unified, the insights that end up being pulled from them are at least somewhat, if not drastically, out of date.

Not only is manual data integration incredibly taxing, but putting the onus of collating copious amounts of data onto employees opens up the risk for all too natural human error

UNCOMMON AD FRAUD TECHNIQUES

Monkey patch exploitation

Monkey-patching is a technique that allows a developer to alter the behaviour of a piece of code at runtime, without changing the underlying code. Whilst useful for developers by allowing them to debug code, this technique can also be used to commit fraud: malicious actors can exploit monkey-patching to alter the original code, disrupting the developers' intentions for how the program is run.

An example of this occurred in the "DefPackage" collection of mobile apps White Ops examined in a recent investigation. Monkey-patching was used to modify the advertising SDK to bombard the user with ads from the moment they unlocked their phone. Those behind the operation had also used the technique to alter functionalities, such as the back button, to prevent users from exiting these imposing ads.

Monkey-patching can be used on code libraries, methods, or even variables, meaning that fraudsters are able to exploit a technique that is intended to benefit developers across various areas of the back-end.

SSAI spoofing

While server-side ad insertion (SSAI) offers apps a sophisticated way of serving ads without disrupting user experience, the technique is nonetheless only in its infancy. And as is the case with all new technologies, malicious actors are proving more than capable of finding and exploiting the gaps in the system.

Fraudsters can exploit SSAI by sending ad requests from data centres for non-existent ad slots. The fraudsters use their server to call the reporting APIs to indicate that the ad has been "shown", when in reality, it is never presented to a real user. This "SSAI spoofing" thus enables bad actors to pocket the money intended to ensure that ads reach real consumers. Often, the only information advertisers receive in an SSAI environment is limited to the "device" user-agent and IP address (sent via the untrusted server), sent in the "X-Device-User-Agent" and "X-Device-IP" HTTP headers. Whilst forging this data is trivial for a bad actor running a malicious SSAI server, doing so convincingly en-masse is rather tricky. Therefore, successfully imitating these headers is really quite a sophisticated type of bot attack.

SSAI spoofing was the core method implemented by the fraudsters behind the ICEBUCKET operation (the widest reaching CTV fraud operation ever), which was uncovered by White Ops in April 2020. ICEBUCKET used custom code to pose as a legitimate SSAI provider for numerous apps and devices and then put together requests for ads to be inserted into video content. However, neither the viewers these ads were purportedly being shown to, nor the CTV and mobile devices they were supposedly being shown on actually existed, they were entirely fabricated — the fraudsters fell into the trap of using outdated device models that are no longer used en masse, and IP addresses that were clearly, but subtly, algorithmically generated.

Exploiting ASNs

Our uncovering of the ICEBUCKET operation also highlighted how cybercriminals can misuse Autonomous System Numbers (ASNs) to carry out their misdeeds. ASNs are unique numbers used to identify each of the autonomous systems that form the back-end of the internet. Our investigation into ICEBUCKET found that the ad requests sent by the operation came from a small set of ASNs. Whilst we can't be certain of why those behind the scheme chose to use these particular ASNs, we suspect that they were appealing to the fraudsters because:

01.

There was a lower degree of regulation from network operators around malicious activity within their data centres;

02.

these ASNs were accompanied by cheap Virtual Private Server (VPS) services.

All of these factors likely made those behind ICEBUCKET confident that they would be able to operate without detection. Regardless of the specific reasoning, however, what is most apparent is that bad actors are taking advantage of unregulated and unsecured ASNs to carry out their fraudulent activities.

FRAUDSTERS ARE CONSTANTLY EVOLVING

That mental image you have of a hacker or fraudster? It's not true, not anymore anyway. The groups behind the fraud operations that White Ops sees are incredibly professional — they have two-week development sprints, their employees get benefits... they know what they're doing and how to do it efficiently.

Our recent BeautyFraud investigation provided a great insight into how the industry currently tackles in-app fraud. Our findings indicated that, on average, a new app was launched every 11 days on the play store, and then removed after 17 days. The speed with which the Play Store removed these apps is encouraging, but the figures paint a picture of an ongoing game of catch-up. What's more troubling is that these apps managed to accrue a substantial number of users in their short life-spans: the apps analysed acquired 565,833 users on average.

What's most likely is that the fraudsters would make a note of how quickly something was pulled down, change their tactics, and re-upload under a different name. That would give them insight into what tactics trigger the app store's review and removal, and they can continue to tweak from there to find a version that lasts as long as possible—making the fraudsters as much money as possible — before removal.

Whilst parts of the code that put the apps in the fraudsters' control remained, enough had been removed to render the fraud activity inactive. It's not exactly clear why they chose to do this — we speculate that it may have been an attempt to find out whether the code caused their apps to be detected and removed, or to preserve the apps for use at a later date. Whether or not either of these prove to be the case, the finding reiterates how fraudsters are constantly working to develop new ways to carry out their nefarious behaviour. That's why it's vital that the industry does not just remain vigilant, but actively anticipates where and how fraud may emerge.

COMBATTING AD FRAUD

Fraudsters aren't going to disappear any time soon, but there are several steps that we as an industry can take to neutralise them. At White Ops, we monitor pre-bid traffic for threats, which allows us to automatically block fraudulent traffic and ensure that money does not go into the pockets of fraudsters. But industry partners can protect themselves by making use of initiatives that help protect them: ads.txt, app-ads.txt, sellers.json, and ads.cert. Proper and widespread use of these will help create a transparent supply chain for advertisers and their partners.

Brands can defend themselves from criminal operations by using marketing fraud insights to optimise their campaigns and acquisition strategies, which will help to spot and remove specific platforms of inventory sources that drive fraudulent traffic and doubling down on clean ones.

Brands can also use these insights to remove bot traffic from their data management and attribution platforms. Doing so makes for cleaner data, which in turn drives better analysis: brands are able to better understand customers when data isn't clouded by bot activity. In a matter of weeks, all these efforts can result in a sizable uptick in a brand's conversion rates and eliminate wasted spend on non-human traffic.

Therefore, not only can addressing fraud produce direct results by preventing spend from going to waste and being diverted into criminal enterprises, but it can provide a host of indirect benefits. Closely examining fraud on a real-time basis gives marketers the ability to optimise conversions, improve data hygiene, reduce costs and drive real human engagement. Marketing fraud is preventing true marketing effectiveness, and the remedy for digital marketers is to keep it human.

It's vital that the industry does not just remain vigilant, but actively anticipates where and how fraud may emerge

PART 2: FRAUD IN EMERGING SECTORS

Fraudsters have one specific pattern of behavior: they follow the money and they work their hardest to take advantage of perceived gaps in the ecosystem, especially during times of change or turmoil.

The excitement that advertisers, consumers, and technology partners experience when a new market emerges can turn quickly into concern when campaigns don't perform as hoped... especially when it becomes clear that fraudsters are clever and got there first.

In this section, we'll explore fraud models and threats to new and emerging sectors in digital advertising.

Connected TV (CTV)

Connected TV (CTV) provides massive opportunities for streaming services and brands to engage with consumers through compelling content and advertising. The cord-cutter movement continues to pick up steam, and new providers are offering streaming services regularly. Research released in April 2020 suggests that the average consumer subscribes to four paid streaming platforms, and the social distancing measures brought on by COVID-19 may have pushed that number even higher since then.

However, where there is growth, there will inevitably be bad actors — there have already been instances of ad fraud in CTV, and unless a concerted effort is made to protect the sector, there will likely be many more. The CTV ecosystem and brands must work hand-in-hand within a collectively protected advertising supply chain to ensure that fraud is recognised, addressed, and neutralised as quickly as possible.

A collective approach needs to take precedence, as fraud schemes rarely target a single publisher at a time. The ICEBUCKET bot operation detailed in section one counterfeited more than 300 different publishers at the time of its peak in January 2020, and impersonated more than 2 million individuals, accounting for nearly 28% of total programmatic CTV traffic at the time. With the threat yet to be wholly eliminated, this is not something the industry can take a back seat on.

Where there is growth, there will inevitably be bad actors — there have already been instances of ad fraud in CTV

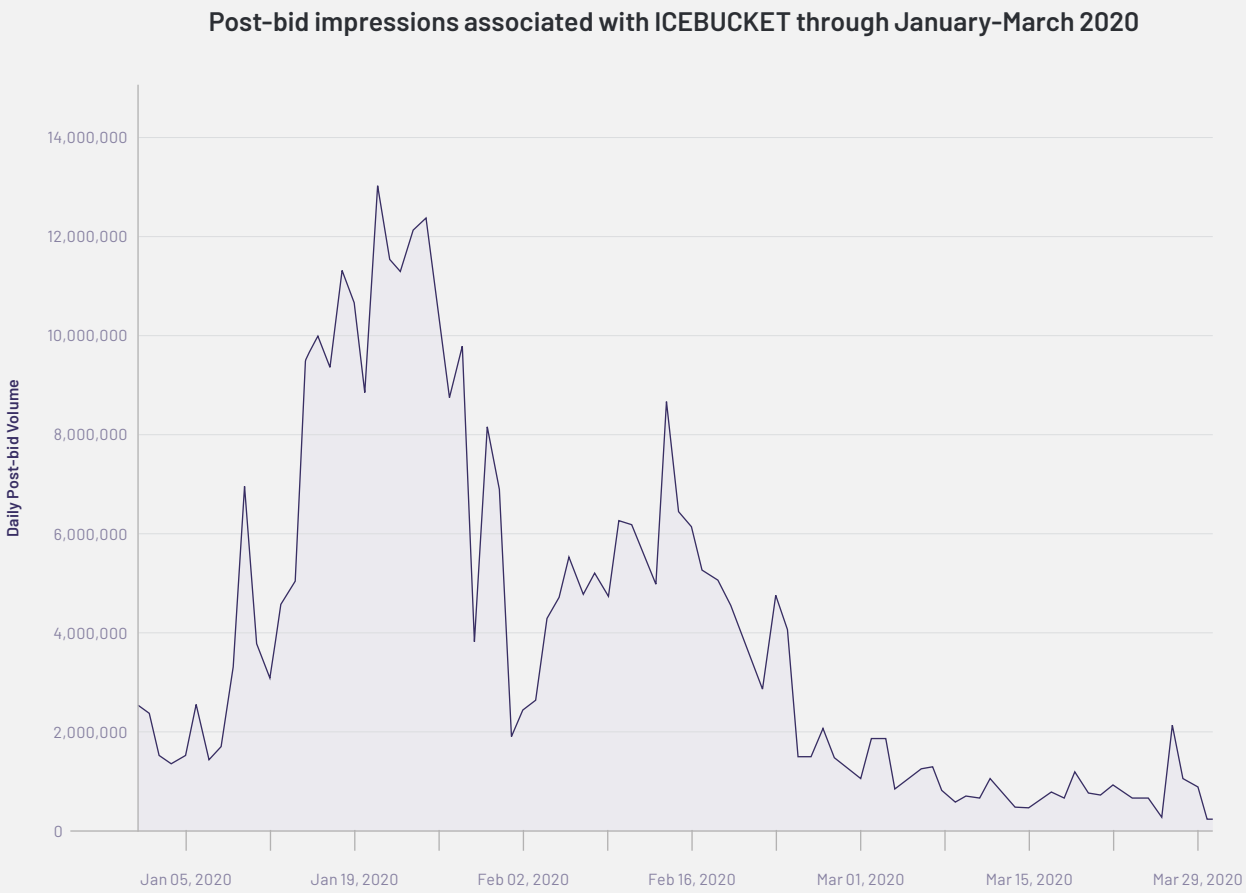

Buying inventory through unprotected channels leaves marketers particularly vulnerable to ad fraud, whereas protected channels which host direct relationships, trust, and full transparency are better equipped to tackle it. Working together through a collectively protected supply chain will ensure the ecosystem realises the full benefits of creating a great CTV customer experience without being marred by ad fraud.

Working together through a collectively protected supply chain will ensure the ecosystem realises the full benefits of creating a great CTV customer experience

In-app

Mobile ad fraud has been a focal target for malicious actors in previous years. It's no surprise that marketers and brands have caught on to this particular threat and are working to mitigate its impacts, resulting in a decrease in fraud in recent months.

But while this decline should be celebrated, publishers must not become complacent. We have no reason to believe that fraudsters will simply give up on their enterprises when they're disrupted: rather, they will adapt their methods to circumvent the industry's latest defences (such as by pivoting towards cost-per-action (CPA) in highly-scaled verticals such as gaming), and will likely continue to launch complex bot attacks. Furthermore, with mobile usage proliferating across Asia-Pacific, the Middle East, and Africa, it is critical that app publishers in these regions keep on top of the weaknesses in their software. Where faults go unnoticed, in-app fraud will resurge.

Audio

Like CTV, programmatic advertising opportunity offered by audio is unfortunately tempered by the risk of fraud. We have already seen bad actors intrude into the audio sphere, with a number of fraudulent podcasts posing as popular programs swamping Spotify's subsidiary platform, Anchor.fm. And as audio formats such as digital radio, podcasts, and music streaming services begin to use programmatic audio advertising more and more, the risk of ad fraud is only likely to increase.

As a more intimate medium than TV or video, audience backlash against audio ads tends to be stronger than that against other channels. Therefore, excessive or out-of-context units driven by fraudsters may prove more damaging to both brand and audio publishers if left unfettered. Given audience aversion to audio advertising, as well as an increased consciousness around the collection and use of their data, we may also see a rise in sophisticated bots and other attempts by fraudsters to fill the space left by consumer data opt-outs with high-value targets.

DOOH

While fraudulent techniques involving non-human traffic are negated by the fact that it is a one-to-many (rather than one-to-one) medium, digital out-of-home (DOOH) is becoming increasingly vulnerable to fraud. This is due both to the increasing digitisation of inventory and growing prevalence of programmatic out-of-home (OOH trading).

With a greater understanding of the potential benefits and dangers of DOOH yet to be achieved, the industry must be flexible yet vigilant in its approach to fraud within the channel. The entire supply-chain would also benefit from greater education on the medium in order to reconcile where DOOH will mimic other sources of programmatic inventory, where it will differ, and where the weak spots that may make it susceptible to bad actors lie.

PART 3: HOW INDUSTRY PLAYERS ARE ADDRESSING AD FRAUD

We've detailed some of the more common and relatively new tactics used by fraudsters to scam advertisers out of their marketing budgets, and have outlined the approach we believe the industry should take to disarm them. But how are members of our community currently addressing the risk of being targeted by bad actors? We spoke to Vodafone's Marcel Zielke and Mindshare's Jan Montwill to learn how their companies are tackling ad fraud.

MARCEL ZIELKE, VBAT, ZV, CHANNEL COORDINATOR, DIGITAL & TECH, VODAFONE GMBH

"Advertising fraud is one of today's most expensive marketing problems, especially for brands. By working with ad verification providers, we can improve the effectiveness of our online advertising, giving us clarity and confidence in our digital investments. At Vodafone, we monitor developments in both open web and social platform ecosystems very carefully; it's integral that we understand where we can receive the best standards in the digital market, as well as what technological safeguards are available to safeguard against fraud, alongside blocking hate speech or other unwanted content. In addition to that, we have a number of initiatives in place, including working closely with all partners and the continuous evaluation of the supply chain, to enhance transparency.

With the implementation of stricter data privacy regulations, such as GDPR, and the phasing out of third-party cookies, the future of audience targeting remains somewhat uncertain, potentially creating an environment in which fraudsters can thrive. However, we are actively reconsidering the potential of contextual targeting, which has become more relevant for our brand and campaigns. Implemented in an intelligent coexistence with our first-party data strategy, we're confident that we're prepared to tackle whatever fraudsters try to throw at us during this period of industry-wide upheaval and transition."

JAN MONTWILL DIGITAL DIRECTOR, MINDSHARE SWEDEN

"In the Nordics, we have, until now, largely been spared from serious ad fraud. The levels of ad fraud are relatively low here compared to other, larger markets when looking at the latest (eMarketer) stats, but that doesn't mean we should take ad fraud any less seriously than operators in more badly affected markets.

This paper shows that ad fraud is evolving across all channels (including channels that have previously been less affected, such as audio and DOOH), reiterating the importance that our collective response should also be evolving. We as an industry must collaborate to find common solutions to tackle this issue: close co-operation with anti-ad fraud vendors is paramount (obviously), but inter-agency and cross-market collaboration must also be a top priority for everyone within the "madtech" industry, as no one sector is immune to ad fraud. In essence, regardless of whether we work in markets where the current ad fraud level is low or high, sharing insights is key to combating ad fraud. Sharing is caring!"

ABOUT WHITEOPS

White Ops is a cybersecurity company that collectively protects global enterprises and internet platforms from digital fraud and abuse. We verify the humanity of more than 10 trillion interactions per week to enable our customers to avoid risk to their data, reputation, compliance, bottom line and customer experience as they grow their digital business.

To learn more and 'Know Who is Real', visit www.keepithuman.com

ABOUT EXCHANGEWIRE

ExchangeWire tracks global data-driven and programmatic advertising, media buying trends, and the ad tech and mar tech sectors. Delving deep into the business of automated media trading and the technology that underpins it across multi-channels (online display, video, mobile and social), the site aims to keep readers up to data on all the latest news and developments.

ExchangeWire provides opinion and analysis on the following sector companies: specialist media buyers, ad traders, ad networks, media agencies, publishers, data exchanges, ad exchanges and specialist ad tech providers in the video, mobile and online display markets.