adform

What Is Ad Fraud and How Can It Be Prevented?

adform

Ad fraud is a term that is familiar to most people in the digital advertising industry. However, due to its complexity it can be difficult to get insight into the full impact that ad fraud has on marketers' campaign results and the industry as a whole. Marketers across the industry regularly express a frustration with digesting the different types of ad fraud and mapping what steps can be taken to help prevent it. To help navigate the topic, in this white paper we will look at:

- The Origin of Ad Fraud
- Three Main Categories of Ad Fraud
- Solutions to Address the Common Types of Ad Fraud
- The Development of Ad Fraud in Emerging Channels

History: Where Did Ad Fraud Originate and Why Is It so Challenging to Address?

When digital advertising was in its infancy, it was treated much like traditional print advertisement. When an advertiser decided to place an ad on a website that their customers would visit, they needed to call the publisher that owned a given website and set up a case-by-case agreement to determine the details of the transaction. This would include when and where the ad would be shown, for how long, and the manual transfer of the creative assets. This led to personal relationships between advertisers and publishers and relied heavily on trust paired with a very transparent insight into where an ad should be displayed.

However, as the internet and digital advertising expanded, it became impractical and cost prohibitive for buyers to maintain these relationships. They found that in order to reach their customers most effectively, they needed to place even more ads on an increasing number of websites. At the same time the opportunity arose to not only reach a static audience, but to identify and then speak to highly engaged users in a personalized way. Having a phone call or meeting with each publisher was time consuming and complex. Ad servers were developed to automate this auction process and make the bids on the behalf of the advertiser. This resulted in transactions becoming less personal. More complex and abstract relationships also made maintaining trust and transparency more challenging, while at the same time it became apparent that advertisers needed better auditing tools that stretched beyond black box setups and unsubstantiated statements from publishers.

As the complexity of the market grew, and significant investment was pivoted to digital channels, a highly lucrative door opened for criminals to leverage ad fraud to develop a sophisticated money-making scheme. These schemes have continued to evolve from relatively simple fake ad clicks to more sophisticated techniques. The scale of the web, the complexity of the technologies and layered partnerships that helped enable displaying ads to the right people at the right time also inadvertently created added opportunities for exploitation by bad actors.

SAMPLE FRAUD NETWORK

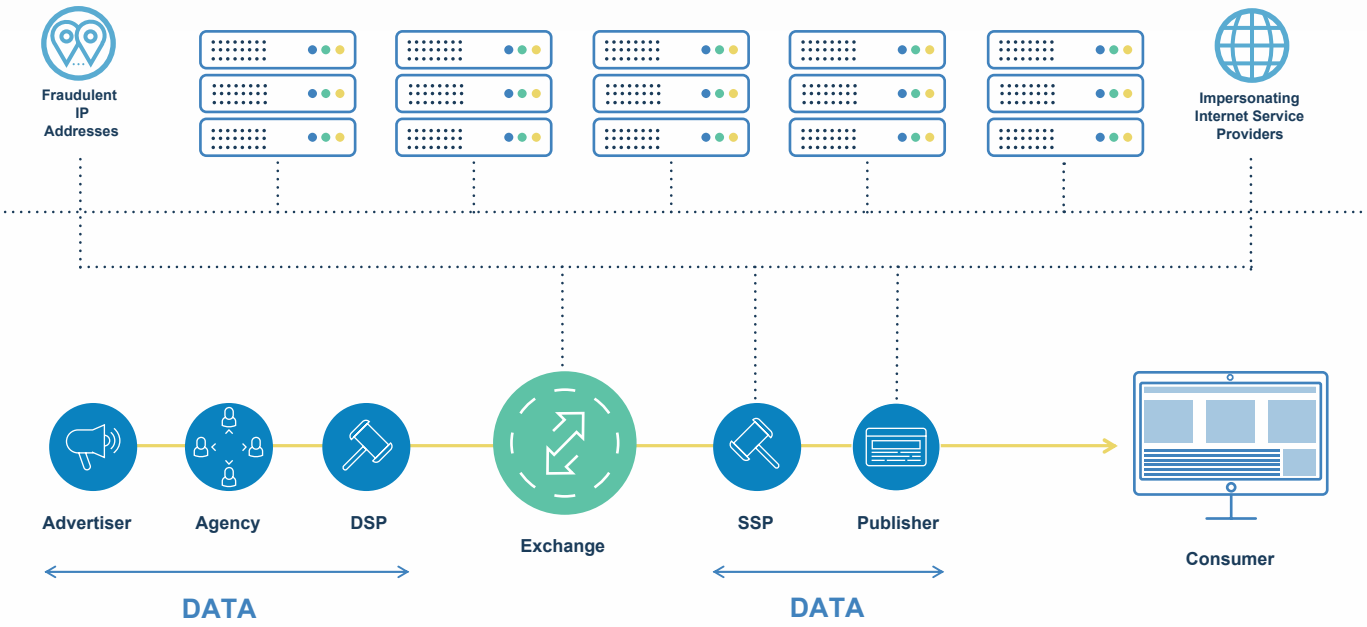

adform

Size and Scope of Ad Fraud

It is difficult to get an accurate view of precisely how large the impact of ad fraud is on the digital advertising industry because even what constitutes digital advertising is constantly changing and can be defined broadly or narrowly depending on case. But the evidence indicates that "fraud attempts amount to 20 to 35 percent of all ad impressions throughout the year" according to White Ops' Bot Baseline Report for 2018-2019. In Adform's Integrated Advertising Platform benchmarks, with the proactive anti-fraud solution Bearskin, fraud is seen in 2-10% of impressions depending on the type, format and location of the ad.

The reason for the difficulty in exactness is because many elements of ad fraud share traits with money laundering schemes seen in traditional brick and mortar lines of business that mix legitimate and illegitimate behavior. Reflecting the complex nature of the issue, primary industry estimates currently vary drastically on monetary impact. Juniper Research estimates that there will be approximately $42 Billion lost across online, mobile and in-app channels in 2019. While White Ops and the Association for Nation Advertisers (ANA) estimate that this figure is closer to $5.8 Billion. Though it stems from many different obstacles, one primary reason for the discrepancy is the fact that the estimate of the total spend is based on publishers' revenue and then calculated backwards to the advertisers' spend. These are self-reported figures and not all publishers disclose this information. In addition, the portion of the ad spend going to fraudsters is not included in this total since they are obviously not reporting their income to the authorities. To complicate the calculations further, the price of ad impressions varies radically based on the format, location and audience so these figures are based on an average making it incredibly difficult to calculate the total monetary value of the industry. As an example, when Adform discovered Hyphbot in 2017, it was estimated that this botnet alone was responsible for over $1.2 million in media wastage per day.

No matter what the actual figure might be, ad fraud is a problem that needs the attention of all players in the ad tech ecosystem. Advertisers are losing money when they pay for ads that are shown to bots and not actual consumers, which impacts the validity of their campaign results. Legitimate ad tech vendors are also impacted by fraud as they spend millions each year on bandwidth, lost bids, fending off attacks, etc. This becomes a major added cost of doing business to good ad tech vendors. Publishers are also losing money because the fraudsters are causing CPM prices to be artificially lowered by crowding ad exchanges with illegitimate ad placements. Marketers are under a lot of pressure to get the best results for their campaign with as little money as possible, so the cheaper prices of these placements are appealing, especially because the marketers believe that they are bidding on impressions from reputable publishers. If the marketers do not carefully scrutinize the results of these placements, it could be easy to assume that the reason for an underperforming ad campaign lies elsewhere and not because the ads are not being seen by actual consumers. By eliminating or at least reducing fraud, advertisers would get a higher conversion rate on campaigns and they would, in turn, be more willing to pay a higher CPM for the legitimate impressions.

The never-ending challenge of fighting ad fraud is that it is difficult to determine what is fraud and what is human activity. As fraud detection technology advances, the tactics used by the fraudsters become more sophisticated, mimicking the clicks and navigation of actual humans which makes detection even more difficult. This creates a continual game of cat and mouse between the industry and the fraudsters.

The digital advertising landscape is also changing and new channels are becoming attractive to fraudsters like CTV and in-app mobile ads where the technological infrastructure is more fragmented leaving gaps that fraudsters can take advantage of in different ways.

A simplified way to look at it would be to break down ad fraud into three general categories: Fake Traffic, Fake Supply and Fake Data with a variety of techniques used for each.

Problem: Fake Traffic

Fake traffic means that fraudsters are falsely inflating the number of visits to websites and website engagements which increase impressions. This can be done in a variety of sneaky ways including:

Human traffic impersonation which can include fraud tactics such as Bots or Click Farms along with other methods of impersonating legitimate users. **Bots**, which are programs that mimic human behavior while interacting with other programs. They are designed to perform tasks like clicks, page loads and video plays of ads in order to trick the publisher into thinking that an actual user has interacted with the ad. **Click Farms**, on the other hand, are actually a group of humans who are paid to interact with advertisements at rapid rates on multiple devices driving up the traffic.

Invalid human activity is when actual human activity is taking place, but the user doesn't realize they are viewing an ad or is not intentionally clicking on a link. This can be done with **Ad Stacking**, which is when fraudsters place multiple ads on top of each other in a single placement, with only the top ad being viewable. The user is unaware that they are "viewing" the other ads in the stack and the advertiser pays for an impression that was hidden. **Pixel Stuffing** is another type of invalid human activity where an ad is placed in a tiny format that cannot be seen by an actual human, but that is counted as a view and the advertiser pays for the impression.

Hijacking is when a user's device (browser, phone, app) is modified to perform a task like loading a page or clicking a link that will create ad requests. In these cases, the user has no control over these activities and they are done without their consent. Hijacking could also include a type of fraud involving legitimate users, but where the device is hijacked to commit fraud during lapses in usage or in the background leaving the user unaware and unable to stop it.

Solution: Advanced Technology

Technology is being continually developed and adapted that can detect and understand these different fraud techniques in order to avoid delivering ads to fake traffic.

The technology monitors and reacts to fraud at different points during the transaction process. It can be used during the pre-bid phase to identify suspicious domains, IP addresses of the User or even Cookie ID and refuse to bid in those auctions. Or, if the technology does not detect a threat until after the ad has been served, it can report and not register the fake impression, logging the fraud to be avoided in future bids.

Adform's Bearskin is an excellent example of technology that addresses this type of ad fraud.

Bearskin is our proprietary built-in fraud detection solution which adds an extra layer of defense and protects our programmatic traffic from fraud, bots, and bad actors.

Bearskin, part of Adform's Odin AI, is a proactive anti-fraud solution. It is based on an algorithm that has evaluated and learned from billions of transactions and data points drawn from across the entire Adform Integrated Advertising Platform, not just the DSP, which gives it greater insight into human behaviors.

Problem: Fake Supply

This kind of ad fraud occurs when advertisers bid on illegitimate placements thinking that they are authentic. Again there are a variety of tactics that fraudsters employ when creating this type of activity, the list below is just a sample of some of the more common types of Fake Supply.

Domain Spoofing is one type of fake supply that has many versions. In its simplest form this involves URL Substitution where a fraudster substitutes a fake URL at the time of the bid that will result in the ad being served in a different location than was bid on. Another version of this is when a publisher owns several domains of varying quality or reputability and misrepresents lower quality traffic during the auction by masking them all under one domain.

Cross-Domain Embedding involves two sites being connected by a fraudster through an iFrame. One site has high traffic but low quality content that would generally have difficulty attracting demand from advertisers. The second site has higher quality or safe content but low traffic. Advertisers believe they are bidding on ads on the safe site and the reporting will show that it has been served in that location. However, the fraudster is able to link the location to the site with the higher traffic through the iFrame in order to trick it into appearing as if it has been viewed by users.

Custom Browsers are a more complex version of domain spoofing. These are used for creating websites that aren't reachable by normal human users. The fraudsters set up servers and put fake websites on them so that they can then send an ad request falsely representing another legitimate site. This tricks advertisers and ad exchanges into thinking the inventory is good and they choose to bid on the placement.

Ad Injection, is when fraudsters hack another web page to insert ads without the consent of the publisher. These ads can be visible to the user or they can be hidden behind other content.

adform

Solution: Adoption of Industry Standards

In May 2017, the Interactive Advertising Bureau (IAB) Tech Lab launched ads.txt, which is a text file that publishers can place on their sites that gives ad buyers a list of the authorized vendors that are allowed sell inventory on the site. This initiative is an attempt to address the problem of domain spoofing giving an added layer of transparency and control across the industry. Since its introduction, it has had an impressive adoption rate with nearly 80% of publishers that offer video inventory using programmatic participating. Studies have been done to determine how much of the inventory on ad exchanges that seems to belong to a participant publisher was legitimate. When the ads.txt filers were removed from ad buys, the study found that nearly 72% of video ad spend went to unauthorized programmatic platforms. In the case of 3ve, the largest fraud network known to date, experts agree that if ads.txt was in place it could have prevented up to 80% - or around $24 million- of their activities.

In March 2019, IAB Tech Lab released the next iteration, called app-ads.txt which is created to support mobile app, OTT, or any other app inventory. Adform's DSP is actively using this level of protection to filter unauthorized sellers from their platform.

In order to address the prevention of ad injection, Ads.cert: Signed Bid Requests was introduced in November 2018 as part of IAB Tech Lab's OpenRTB 3.0 updates. Ads.cert is a public key system that lets buyers verify if any of the details from the original ad request were altered at any point in the supply path. This gives buyers the ability to check that they are getting what they paid for and flags changes in placement or device type that may have been tampered with.

Using ads.txt and ads.cert together, the buyer is able to check that the seller has the rights to sell the inventory they want to buy as well as making sure that the inventory is accurately represented during bidding and delivery.

adform

Problem: Fake Data

This is when the fraudsters target the data as it is returning to the buyer and hack the measurement in order to show more impressions with high viewability or website visits to fake good performance. While the difference in how you access and explore the data means that this type of fraud is less readily visible, it is worth investing the time to digest the data and familiarize yourself with its patterns. As a result, you'll be able to identify and flag suspicious data and unusual outliers more rapidly.

Solution: Advertiser Responsibility

This is the piece where the advertisers can take ownership of their data and raise a red flag when something seems off. In order to do this, "marketers" need to first be aware of ad fraud and the impact it can have on their campaigns. In order to know if the data is being tampered with, advertisers can start by defining their KPIs, measuring the results, analyzing the campaign data and then taking action if something looks strange.

By setting up these steps and monitoring the campaign on a daily basis, the marketing team is able to understand how the campaign is actually performing. This allows them to recognize when the pattern changes and then look into where the money is going. Being able to monitor your results at a granular level and have ownership over your data is vital to having a clear picture of performance. Taking an active role in examining these results can help save advertisers valuable resources and prevent their hard work from being wasted.

As Dr. Augustine Fou has said,

> "If they see strange things in the analytics, don't ignore it; ask questions. There has to be a reason for strange things like 100% bounce rates, 0% bounce rates, perfectly consistent pages per session across dozens of referring sites, etc. It's probably not from real humans seeing your ad and visiting your sites. "

Some important steps to remember when addressing Fake Data:

- Define KPIs for your campaigns
- Measure
- Analyze the Data
- Take Action

Developments in Fraud

Ad fraud is a problem that is contastly adapting and changing as the ad tech industry learns to identify the latest tricks and becomes more capable of shutting them down. Recently, there has been a shift where fraudsters are starting to attack other channels beyond online display advertising. Apps are particularly vulnerable because of the fragmentation of the technology. This leaves openings for fraudsters to enter the ecosystem and redirect the funds to themselves. For example, one tactic fraudsters are using is to purchase cheap in-app display inventory and fill it with multiple video players behind innocuous fake branded display ads.This allows the fraudster to sell the more expensive video ads at a premium keeping the difference in revenue from these ads for themselves. These video ads will be hidden and never seen by a customer.

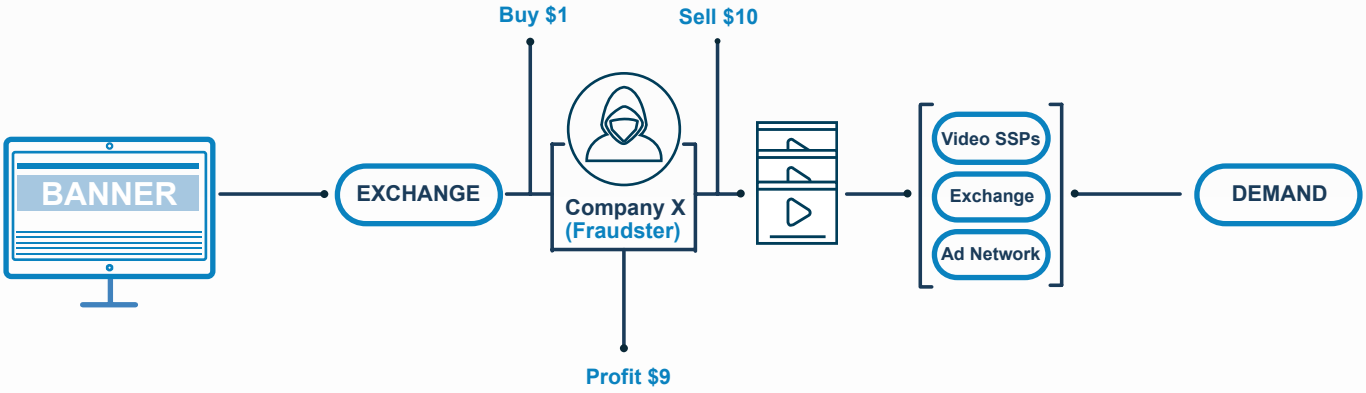

Another newer area being targeted by fraudsters are audio streaming channels. Stream Farms (similar to Click Farms) are being built to play songs and videos on a continuous loop in order to boost the total play count. Video is especially venerable because of the high premium paid for video ads. Social media is also seeing a rise in fraud in the form of influencer bots as well as actual influencers who are engaging in fraudulent behaviors such as buying followers, likes, shares, comments and other engagement activities done by Bots instead of real people.

Another recent example of how the fraud problem is constantly evolving and shifting is the most recent development in faking consent strings. A consent string is “a series of numbers added to an ad bid request, which identifies the consent status of an ad tech vendor. That means whether or not they have a user’s consent to use their data in order to serve them personalized advertising.” (DigiDay 2019) Without a consent string in place for a vendor, demand for impressions on their site decreases substantially. So fraudsters who have set up these fake domains need a way to make them look legitimate so they are creating fake strings or altering existing consent strings. This makes it appear as if the fraudulent vendors have been given consent by a user, when in fact that user denied or failed to provide permission.

How To Take Action

However, key actors in the ad tech industry are working tirelessly alongside legitimate advertisers and publishers to stamp out fraud wherever it exists. In the face of this massive problem, it could seem daunting to try to protect your company and campaigns againt ad fraud. However, the good news is that even taking a few simple steps has the potential to make a significant impact. You can start by:

Understand fraud and its techniques

Own and leverage full access to your data

Analyze and understand your data

Use credible vendors to help identify fraud

Adhere to industry best practices like ads.txt

adform

Adform provides an integrated Software as a Service platform for the buying, managing and serving of digital advertising. The company's software consists of a Data Management Platform, a Demand Side Platform and an Ad Serving Platform with advanced analytics, reporting and creative tools that drive high impact digital advertising campaigns globally. Founded in Denmark in 2002, Adform services a client portfolio that includes the world's leading agencies, advertisers, consultancies, and publishers.

To learn more visit us at **www.adform.com**

Last updated November 2019