AD FRAUD 101

THE COMPLETE BEGINNER'S GUIDE

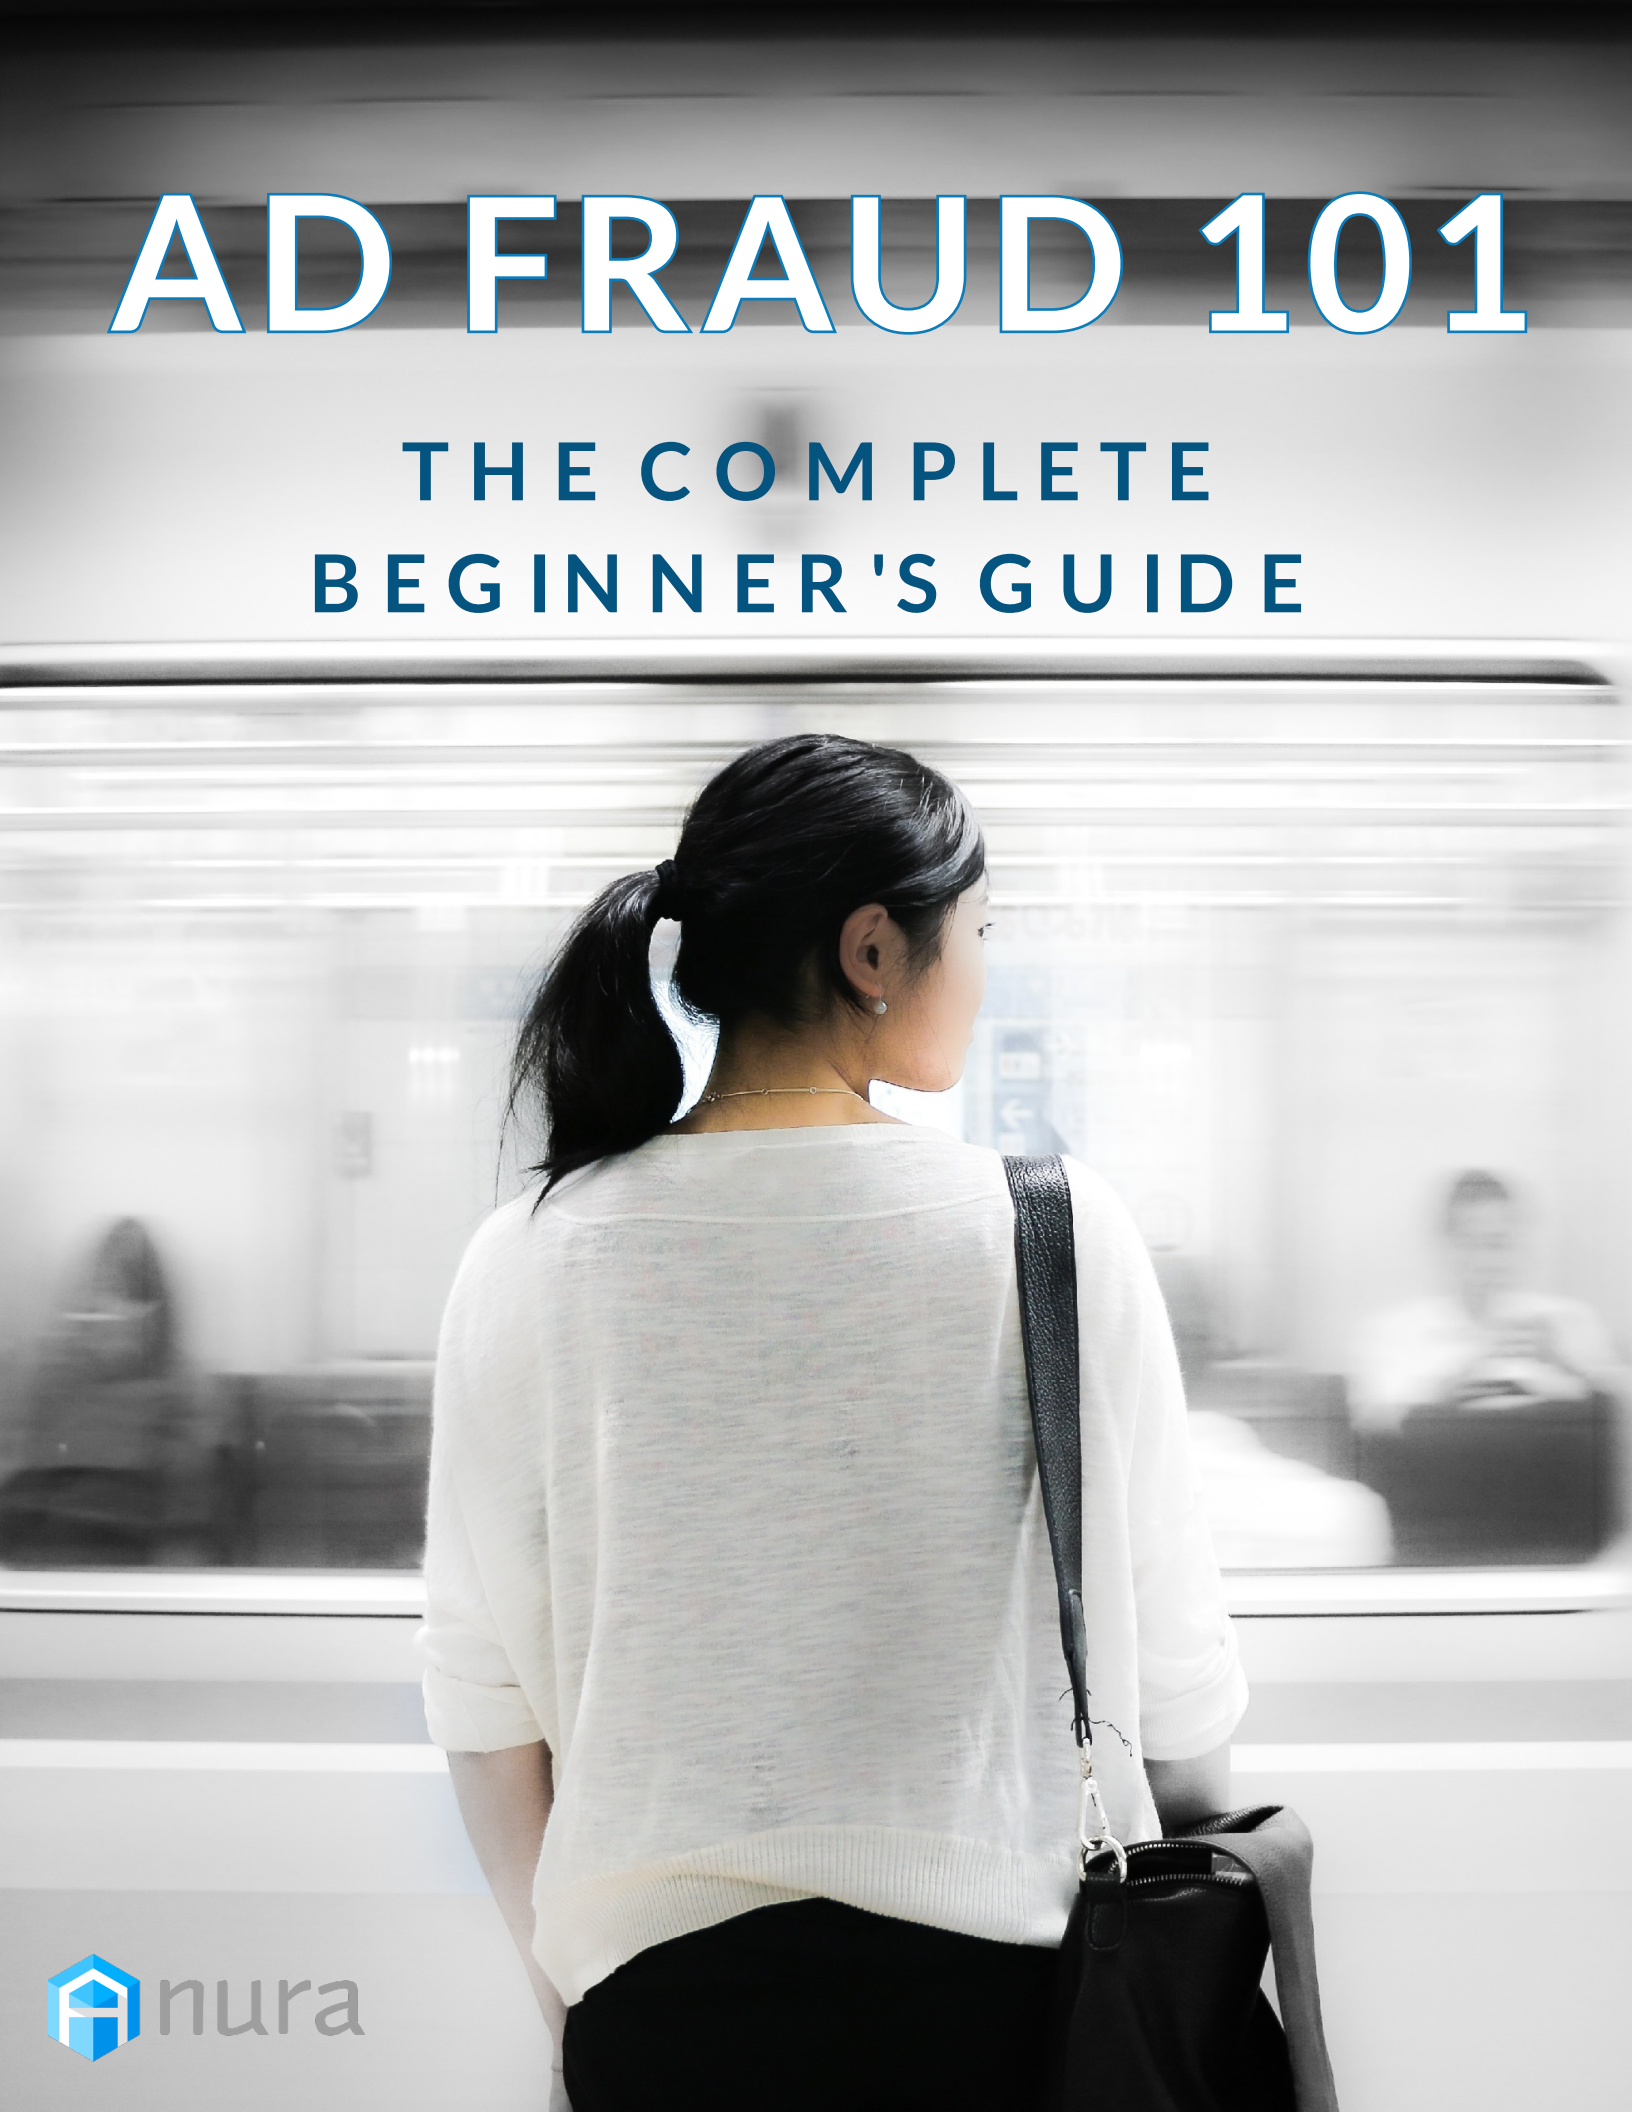

anura

FORWARD

Hey there!

Thank you for downloading our eBook, Ad Fraud 101: The Complete Beginner's Guide. This introductory guide will help walk marketers and advertisers through the complex (and always evolving) world of ad fraud. In the following chapters, you'll find topics ranging from:

- What Is Ad Fraud
- What Are the Signs of Ad Fraud
- Ad Fraud Types
- How to Protect Your Brand

We've also included our robust Ad Fraud Glossary filled with up-to-date ad fraud related terms and definitions to help you cut through the jargon and get the clarity you need. Knowing what you're dealing with is half the battle in fighting ad fraud.

Thanks for reading,

Rich Kahn, CEO and Co-Founder, Anura.io

Rich Kahn is the Co-Founder and CEO of Anura.io, an ad fraud solution that monitors traffic to identify real users versus bots, malware, and human fraud. Anura is the culmination of more than a decade of fraud detection efforts within digital marketing firm eZanga.com, which Rich also co-founded and owns.

Previously, Rich held management roles at Verizon Wireless and Bloomberg, before starting his own internet service provider, First Street Corporation. He co-founded Paid for Surf, an advertising software company, and was the COO of the pay per click advertising network AdOrigin. Rich is considered an industry expert, having over 25 years of global experience with technology, digital advertising, ad fraud management, and elimination.

TABLE OF CONTENTS

WHAT IS AD FRAUD AND HOW IT WORKS

Ad fraud is the practice of viewing, clicking, converting, or generating false interactions with any web asset for the sole purpose of earning money directly or indirectly. And it isn't just bots doing bad stuff on the web. Fraud can also be conducted by humans (e.g. click farms), too.

Fraudsters use a variety of methods to trick unsuspecting advertisers into doing business with them. On the surface, these forms of fraud appear legitimate to an unsuspecting user, but a closer look says otherwise.

For instance, this infographic below shows a human hacker's strategy for using impression fraud. In just six steps, a hacker will successfully execute an ad fraud scheme. Scary how "simple" it is, isn't it?

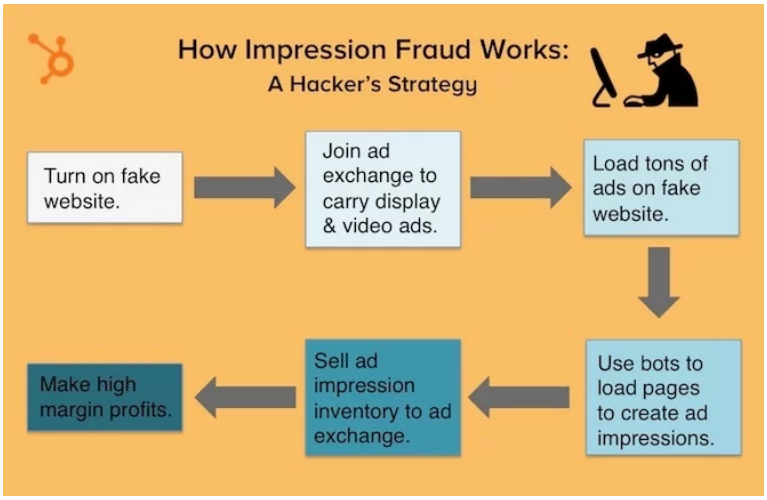

Source: HubSpot

Now that you have a general idea of what is ad fraud, let's review the signs of ad fraud.

AD FRAUD SIGNS

Detecting the signs of ad fraud can be tricky, especially if data analysis isn't your strongest skill. Here are some simple metrics you can track to gauge your click campaign's security.

LOTS OF CLICKS BUT LOW CONVERSIONS

If you're experiencing abnormally high numbers of clicks with little conversions, you might be an ad fraud target. You want people clicking through your ad and completing an action, like filling out a form, subscribing to an email campaign, or making a purchase. In an ideal situation, conversions should increase with click volume; as more people visit, the chances of conversions go up.

Fraudsters aren't trying to convert; their goal is to eat away at your ad budget.

There is also conversion fraud which is harder to recognize. In this scenario, fraudsters "convert" by filling out a form after clicking on your ad. But once they complete the action, the bot or human, exits your site.

HIGH BOUNCE RATES

A bounce is a single-page session on your website. It's what happens if someone arrives at your website's homepage, for instance, but doesn't go any further into your site. You can calculate your overall bounce rate by dividing single-page sessions by all recorded sessions.

High bounces rates aren't usually a good thing, but not all instances are ad fraud-related. Your site may not catch people's attention or have the info they need, so they leave. Some may have accidentally clicked through to your page.

To tell the difference, look for a high click-through rate paired with a high bounce rate. Bots or human fraudsters will click-through your ad, arrive at your landing page, then immediately leave. And you're left paying for bogus clicks.

STRANGE TRAFFIC SOURCES

Think your campaign's been compromised? Dig into your traffic data.

Traffic data can reveal small, critical details that there's an underlying fraud problem. Look for traffic coming from outdated browsers or unpopular devices. Another red flag is multiple clicks coming from the same IP address, or surges of traffic coming from random countries outside of your target parameters.

If you answer yes to any of those questions, you might be a victim of ad fraud. Organized schemes, like click farms, often operate out of foreign countries. Fraudsters might use hosted servers to obscure their locations or operate their bots; traffic generated from these servers could share an IP address.

TYPES OF AD FRAUD

Now that you know the signs of ad fraud, watch out for these types of fraud:

CLICK FRAUD

This common form of fraud occurs through automated clicking by bots or click farms. It's used to create fake impressions, making it look like an ad has drawn more traffic than it really has.

LEAD GENERATION FRAUD

Bot and humans can easily generate ad impressions and fill out forms. They trick advertisers into thinking they've gained a new lead's contact information, but really it's just a fraudster.

TRAFFIC FRAUD

Similar to lead fraud but the publishers are the victims. To increase traffic numbers, some publishers will purchase additional traffic so they can charge more for advertising. However, most purchased traffic comes from unreliable third-party sites which are crawling with bot traffic.

RETARGETING FRAUD

Scammers program bots to act like humans to interact with ad and trigger a retargeting campaign.

SEARCH AD FRAUD

Fraudsters use keyword stuffing to make their phony websites appear higher on search engine results. Advertisers then buy ads on these fake sites with the false hope that their ads will be seen.

AFFILIATE AD FRAUD

Also known as “cookie stuffing.” Fraudsters target the commissions affiliates make by driving consumers to a brand’s website where consumers make a purchase. How it works: scammers track the traffic by using cookies. Then when a tracked consumer makes a purchase, the fraudster drains part or all of the affiliate’s commission.

AD STACKING

Sketchy publishers superimpose a number of ads on top of each other, but only the top ad is visible. Meanwhile all of the ads are registering false impressions.

PIXEL STUFFING

This scheme involves shrinking an ad to a one-by-one pixel size on a page. These ads are invisible to the human eye, but they’ll still register an impression if there is traffic on the page.

DOMAIN SPOOFING

Fraudsters misrepresent their sites as legitimate by making the URL appear to be a trusted domain when it’s not. For instance, they might swap a lowercase “l” for a capital “I” and create a domain that looks like that of a credible company.

AD INJECTION

Malware infects a web page and makes it vulnerable to unwanted ads. Brands like Target and Walmart have fallen victim to ad injection. When an ad is injected onto a site by a fraudster, that ad space is never paid for. And most of the time, the ad will be inappropriate for the site they’re on, or take anyone that engages with them to a questionable site.

AD FRAUD LAWS

In 2018, ad fraud was estimated to have cost advertisers $19 billion. With ad fraud putting such a drain on advertising revenue, it begs the question: why aren't advertisers taking any legal recourse?

LEGAL RESOURCES ARE EXPENSIVE

Legal resources aren't cheap, and they might not be as effective as you'd think. The digital world makes it easy for scammers to commit fraud and then cover their tracks. In fact, your legal team might not even be able to collect enough evidence to hold the attackers accountable. Or, if they can make a case, you'd still be back to square one when another group of fraudsters attack. Either way, you're stuck with a hefty legal bill and still have an ad fraud problem.

FRAUD IS EASY TO COMMIT AND HIDE THE TRACKS

Fraudsters use a variety of methods to conduct fraud: click fraud, pixel stuffing, ad stacking, search ad fraud, domain spoofing, the list goes on and on. For scammers, fraud is easy to commit, and it's even easier for them to hide their tracks. For advertisers, it's virtually impossible to track down the exact perpetrators.

For example, simply masking an IP address can obscure the origins of fraudulent traffic. And who has time to unearth the real IP address? Consequently, many advertisers just write off the loss and move on.

THERE ARE TOO MANY FRAUDSTERS

Ad fraud is a low cost, high reward environment that attracts lots of bad actors. No one has the time or resources to track all of them down.

HOW TO SOLVE THE AD FRAUD PROBLEM

So, with no laws against ad fraud, how do we solve the ad fraud problem?

A.I. FRAUD DETECTION

By using A.I. and machine learning, companies can potentially strengthen their anti-fraud arsenal. They can fight money laundering with real-time analysis of transactions, safeguard cloud storage, reduce false positives, and protect company logos.

MEDIA BUYERS NEED TO ASK QUESTIONS

While they can't stop click fraud, media buyers can help mitigate it by asking questions to avoid worthless traffic. Questions need to include: "what makes a click fraudulent," "where is the traffic coming from," "is there a third-party validation system in place," "what's your traffic filtration score," and "is your traffic filtered in real-time."

USE AN AD FRAUD SOLUTION

There's one metric fraudsters can't beat: validated lead conversions. And to validate leads, brands need to have an ad fraud solution in place. Specifically, it needs to be a solution that detects the real users from malware, bots, and human fraud.

Using an ad fraud solution like Anura saves you time, resources, and money.

AD FRAUD GLOSSARY: TERMS AND DEFINITIONS

There's a lot of terminology in the word of ad fraud. And it's always evolving. Cut through the jargon and fix your fraud problems with our glossary of ad fraud terms and definitions.

AD FRAUD

The practice of viewing, clicking, converting, or generating false interactions with any web asset for the sole purpose of earning money directly or indirectly.

AD INJECTION

This type of fraud occurs when software "injects," or inserts, ads on sites without permission. Usually, internet users unsuspectingly download a browser extension or toolbar that secretly contains injection malware, which then infects their computer and displays unsolicited ads.

AD STACKING

Found on the publisher side, ad stacking is the act of placing ads on top of each other on a page, but only the top ad is visible. However, if a user visits the page, all the ads, even the unseen ones, are charged for an impression.

ADS.TXT

Short for Authorized Digital Sellers, ads.txt is a system of plain text files that publicly disclose which companies are permitted to sell a publisher's inventory. Ads.txt files are hosted at the root level of a publisher's domain, making it easy for programmatic distributors and buyers to access the information.

AFFILIATE FRAUD

To pad their commission payouts, some affiliates use deceptive lead generation practices to drive campaign traffic. This includes buying traffic from unauthorized sources, incentivizing leads, and using bots to click ads and fill forms.

A.I.

No longer the stuff of science fiction, A.I., or artificial intelligence, refers to a computer program's ability to simulate intelligent human actions. A.I. gains its abilities through a developmental technique called machine learning.

BLOCKCHAIN

A blockchain is a type of data structure that identifies and tracks transactions across a network of computers, creating a transparent and secure digital ledger.

BOT

Derived from "robot," bots are automated software programs designed to carry out specific tasks. While there are both good bots and bad bots, most people refer to the malicious kind when talking about these programs.

BOTNET

Usually controlled by a remote source, a botnet is a group of malware-infected computers that work together to perform specific tasks, usually without the computer owner's knowledge.

BRAND SAFETY

This term refers to the methods used by advertisers to ensure an ad doesn't appear next to questionable, potentially brand-damaging content. Brand safety is often subjective; for instance, a family-friendly brand like Disney wouldn't want their ads placed on a site that promotes alcohol. But an edgier brand, such as MTV, might not mind.

CALL FRAUD

Ad fraud extends to phone calls, too. Call fraud happens when fraudsters generate phone calls that pass a billable requirement, like call length. Often, they use unauthorized data or prerecorded messages to complete a call and generate a fake lead.

CAPTCHA

Short for “Completely Automated Public Turing Test to Tell Computers and Humans Apart," a CAPTCHA is a visual or audio-based puzzle designed to block bots from executing certain actions, like filling forms or posting comments.

CLICK BOTS

A type of malicious bot, a click bot is designed to engage with advertising and target cost-per-click (CPC) ad campaigns. Advertisers end up having to pay for phony clicks caused by these bots.

CLICK FARMS

These organized operations consist of large groups of low-paid workers that are hired to click on ads and links in order to benefit a head fraudster. Many click farms are found in foreign countries.

CONNECTED TV (CTV)

Not to be confused with Over-the-Top (OTT), CTV requires a smart device to access content. By using your smart TV or streaming device (e.g. Roku), you consume IP delivered content through an app.

CONTENT FARM

A content farm is a website that hosts a large amount of low-quality content in order to game search engine algorithms. Content farms make sure their content ranks high on search results, which in turn gives the illusion that the site attracts lots of traffic. As a result, advertisers are more inclined to buy space on a seemingly popular content farm.

COOKIE STUFFING

It’s not as delicious as it sounds. Mainly affecting affiliate marketing, cookie stuffing involves planting third-party cookies onto users after they visit a website. However, the cookie they get is from a completely unrelated website. Fraudsters use this “stuffed” cookie to track affected users. If any of those users completes a purchase, the fraudster gets credit for the referral, not the affiliate.

CPA FRAUD

Cost per acquisition (CPA) fraud is usually associated with affiliate fraud. Affiliates get rewarded for promoting a product and sending people to a brand's website, where they might convert in some way. Fraudsters commit CPA fraud when they steal referral credits from affiliates. Types of CPA fraud include practices like cookie stuffing, mentioned above.

CPC FRAUD

Cost per click (CPC) fraud is a broad fraud category that mainly deals with search manipulation. Generally, fraudsters target the most popular search keywords, as those yield the biggest payout. The most common method of committing CPC fraud involves creating phony sites or content farms filled with high ranking keywords. The sites look legitimate, letting fraudsters sell ad space at a premium.

CPL FRAUD

Cost per lead (CPL) fraud is committed by both bots and humans. Here, fraudsters falsely generate leads that will never convert. Some CPL fraud actions include filling out forms and clicking on ads or CTA's in order to trigger a payable event. Retargeting fraud, described below, is a form of CPL fraud.

CPM FRAUD

CPM stands for "cost per thousand," and it describes how much money an advertiser will pay for one thousand impressions. CPM fraud is a blanket term for fraud techniques that target impression counts, such as ad stacking or pixel stuffing.

DOMAIN SPOOFING

Here, fraudsters change the URL of their sites so it looks like a more reputable site (think www.paypal.com vs. www.paypals.com). Fraudsters then sell ad space on their fake site at a discount. Spoofed sites might also host dubious content, like adult themes, which could ultimately damage a brand's reputation should their ads appear on the site.

GDPR

The General Data Protection Regulation is a European Union law that aims to regulate the collection and usage of personal data belonging to a living, identifiable EU citizen. Under GDPR, companies are legally required to protect user data as well as empower users to control, monitor, and delete any personal information they want.

HONEYPOT

A honeypot, or blackhole, is essentially a decoy to block bots from filling out your form fills. By placing a dummy form on your site that's only viewable to bots, you're ensuring bots will be drawn to the invisible form instead of the real one.

HUMAN CLICK FRAUD

Similar to click farms, human click fraud is a general term describing the act of using real humans to perform payable actions, like clicks and form fills, in order to deplete an advertiser’s budget.

INCENTIVISED TRAFFIC

Sometimes, in order to boost site traffic, publishers will offer incentives, like monetary rewards, to people who visit the site. Having incentivised traffic coming to a campaign is not ideal because most of the traffic probably won’t convert.

INTERNET OF THINGS

Often abbreviated as IoT, the Internet of Things refers to the network in which any tangible object with internet connectivity can communicate with other connected devices.

MACHINE LEARNING

Heavily reliant on data analysis, machine learning is a branch of computer science that gives computers the ability to learn and improve on specific tasks without direct programming. Essentially, machine learning is what makes A.I. possible.

MALWARE

Malware is often part of the ad fraud mix. This type of malicious software is generally used to gather sensitive information or disrupt computer operations.

OVER-THE-TOP (OTT)

OTT accesses content 'over the top' of an infrastructure provider. With OTT, you consume film or TV content via the internet, not through a traditional cable or satellite TV service (e.g. Comcast or DirecTV).

PIXEL STUFFING

Operating on the publisher's end, pixel stuffing happens when a bunch of ads are served into a single pixel on a web page. Technically, the ads are viewable, but they're physically unseen by users. Advertisers still have to pay for an impression.

RETARGETING FRAUD

This type of CPL fraud involves bots that mimic human behavior. Fraudsters program bots to act like interested customers. For instance, the bots may click through product pages or abandon a shopping cart full of items. These actions trigger a retargeting campaign, causing advertisers to waste money on these fake customers.

SPOOFED TRAFFIC

In the tech field, "spoofing" means masking one's identity in order to deceive or trick another computer or user. Spoofed traffic is traffic that's been manipulated to disguise its origins. Fraudsters use spoofed traffic to make it seem like real visitors are coming to a site from a variety of locations, when in reality, it's probably coming from a single source.

VERIFICATION STRIPPING

By using bots or malware, fraudsters strip the code that's used to verify impressions. They are then able to mask any signs that domain spoofing or other types of fraud is occurring.

UNAUTHORIZED SOURCES

Advertisers generally want traffic that comes from clean, trustworthy sources. However, to meet traffic demand, sometimes publishers acquire traffic from unauthorized sources that are outside of the agreed upon terms.

CLOSING

Dear Reader,

Ad fraud isn't going anywhere. By being proactive, knowing the signs, and adding a layer of protection, you'll make it harder for scammers to commit fraud.

We hope our introductory guide was able to help you out. If you have additional questions about ad fraud, check out our article resource page at https://blog.anura.io/blog.

Thanks again for reading,

Rich Kahn, CEO and Co-Founder, Anura.io

888-337-0641
pr@anura.io
www.anura.io

Anura