iab.
australia

digital ad fraud handbook

2023

thank you to our contributors

André Candeia Galvão
Strategic Partner Lead
Google

Jess Miles
Country Manager ANZ
Integral Ad Science

Imran Masood
Country Manager ANZ
DoubleVerify

Deanna Galluccio
Head of Commercial Ops
Pedestrian Group

Nathan Farrugia
Senior Technical Ops Specialist
REA Group

Colin Lam
Director of Strategic Partner Development
Index Exchange

Declan Dowd
Senior Ad Tech Manager
News Corp Australia

Rahila Nadir
Senior Platforms Solutions Lead
Yahoo

guest contributors

Dr Augustine Fou
Independent Ad Fraud Researcher
FouAnalytics

Lindsay Bender
Head of Product Marketing
HUMAN

Mathew Ratty
Co-Founder & CEO
Adveritas / TrafficGuard

Contents

Introduction 1

The IAB Australia Standards and Guidelines Council have for some time now wanted to provide some general definitions and guidance on ad fraud, in order to support the local industry on a topic that is critical to both understand and take seriously. So, we're delighted to finally have the time and resources to concentrate on this important subject and provide objective guidance. The intention is to build awareness, education and provide a foundational starting point for both buyers and sellers to mitigate the associated risks.

The issue of ad fraud is not a new one. From some of the well-known early issues with the RMX (Right Media Exchange) in the mid 2000's to the dramatic and open acknowledgment by AppNexus in 2015 that it had a problem and the resources and effort it took to genuinely tackle it (as publicly presented at an event in London which I attended), it's been a known and long-standing problem for our industry. Industry focus has traditionally tended to oscillate between complacency - to versions of fear-mongering generated by enormous and often inferred sums of money being criminally syphoned out of the ecosystem.

Whilst acknowledging that these opposing narratives have somewhat been fueled by misaligned incentives in the past, we are simply assuming that the truth lies somewhere in the middle - and have therefore attempted to remain as objective as possible in this guidance and hopefully it's useful to everyone as a result. Ad Fraud is a well-known issue, which demands relentless scrutiny and we hope that brands can take it seriously and with pan-industry support both buyers and sellers can work collaboratively (and closely with those vendors that meet their needs) to minimise these activities, which ultimately negatively impact us all.

We'll start with some general definitions, before providing more details on the different types of fraud that exist along with some practical best practices and guidance. This topic is a persistent one and the threats and considerations are ever-evolving and very fast moving, so please don't assume that this is an exhaustive list. Overall our hope is that with this information and basic guidance our members can balance requirements between their own resources through diligent processes and practices and any dedicated vendors, depending upon specific requirements.

When defining what ad fraud is, we firstly need to look at invalid traffic (IVT), and then dig a little deeper. IVT is essentially any type of traffic that is not from a real user with genuine interest in the content. This can include accidental clicks caused by intrusive ad implementations, fraudulent clicking by competing advertisers, advertising botnets and much worse.

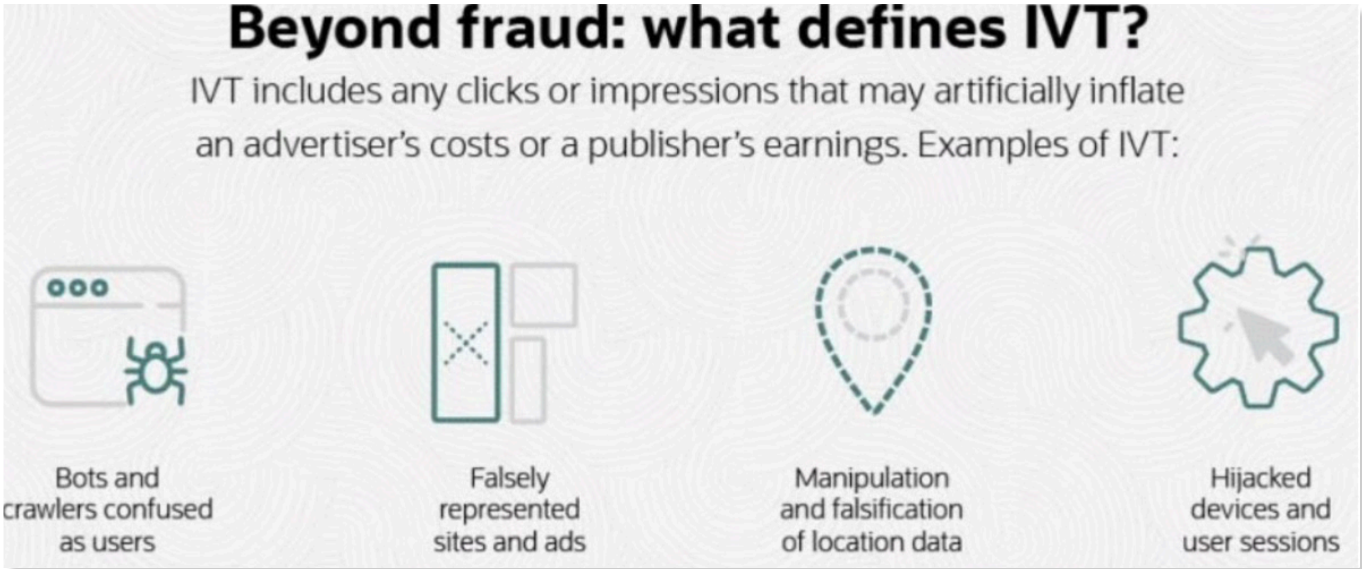

Beyond fraud: what defines IVT?

IVT includes any clicks or impressions that may artificially inflate an advertiser's costs or a publisher's earnings. Examples of IVT:

- Bots and crawlers confused as users
- Falsely represented sites and ads
- Manipulation and falsification of location data
- Hijacked devices and user sessions

image source: Oracle Advertising

From an advertising perspective, the impressions and clicks won't lead to genuine revenue or capture human attention, making them essentially worthless to both publishers and marketers.

So - whilst ad fraud is quite a broad term, it generally refers to both the unintended delivery of advertising to non-human traffic, as well as any intentional activity or action taken by an individual or group of individuals to deceive or manipulate the digital advertising ecosystem for personal gain or financial benefit.

This can include a wide range of fraudulent activities, including fake ad impressions, clicks, conversions, or installs, as well as other deceptive practices such as domain spoofing, click spamming, and bot traffic. Ultimately all these activities result in wastage and lost revenues at best - and at worst directly funds organised crime.

A key part of the problem is the attractiveness to organised criminals of these types of activities, as highlighted in the 'The Business of Hacking' report from Hewlett Packard a few years ago (see below). The digital advertising industry therefore has to remain ever vigilant against bad actors and their relentless initiatives.

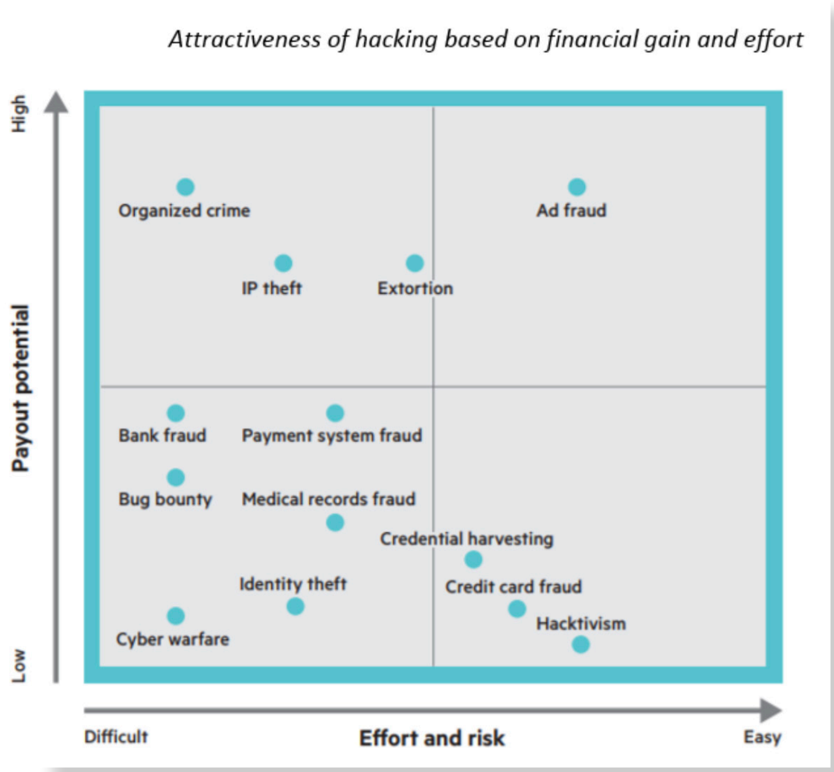

image source: Hewlett Packard

These two different types of IVT can then be segmented further as being either 'General Invalid Traffic' (GIVT) for the more benevolent non-human traffic, or else 'Sophisticated Invalid Traffic' (SIVT) for the more malicious and essentially criminal behaviour.

Both types need to be dealt with and minimised, and the awareness and an acceptance that it's very hard and limiting to attempt to completely eradicate both types entirely from digital activities without seriously compromising results. However, it's not impossible - and simply depends upon a business's requirements, appetite for risk and approach to media quality.

We hope that you find this guidance useful and if you have any feedback at all, or further case studies, please feel free to email me directly at jonas@iabaustralia.com.au

jonas jaanimagi
technology lead
iab australia
jonas@iabaustralia.com.au

definitions of general invalid traffic (givt) & sophisticated invalid traffic (sivt)

2

As mentioned, invalid traffic (IVT) can be split up into 2 different types:

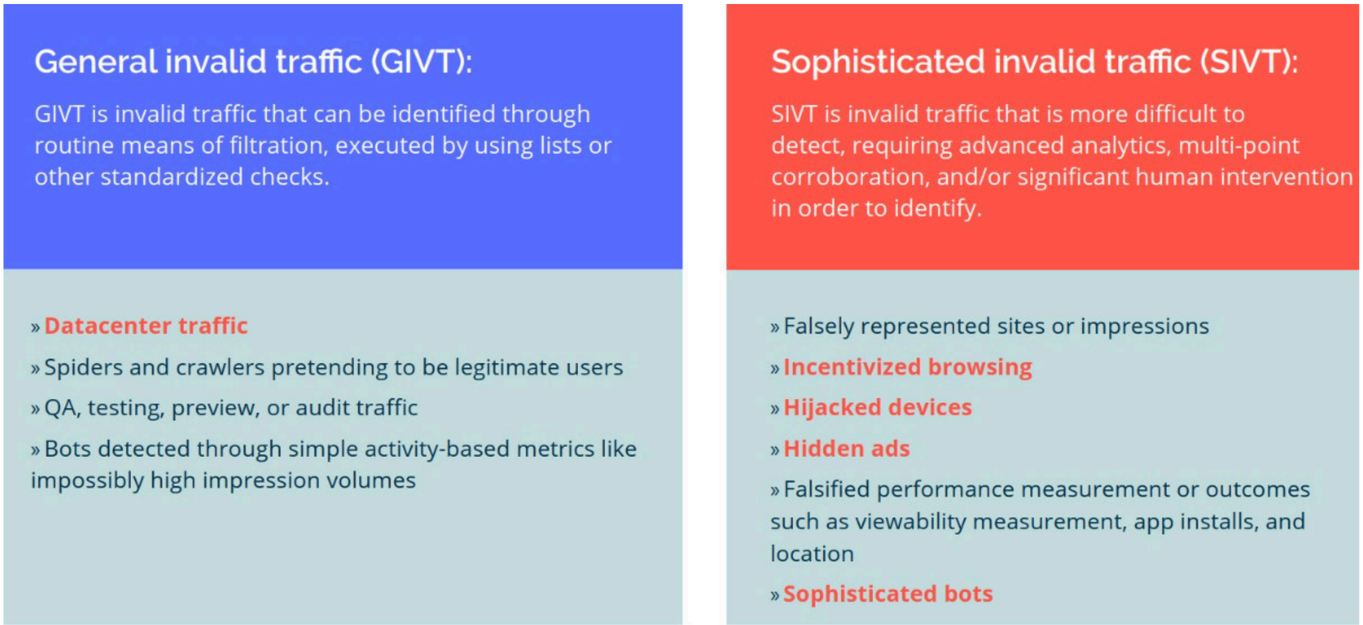

General invalid traffic (GIVT):

GIVT is invalid traffic that can be identified through routine means of filtration, executed by using lists or other standardized checks.

» **Datacenter traffic**

» Spiders and crawlers pretending to be legitimate users

» QA, testing, preview, or audit traffic

» Bots detected through simple activity-based metrics like impossibly high impression volumes

Sophisticated invalid traffic (SIVT):

SIVT is invalid traffic that is more difficult to detect, requiring advanced analytics, multi-point corroboration, and/or significant human intervention in order to identify.

» Falsely represented sites or impressions

» **Incentivized browsing**

» **Hijacked devices**

» **Hidden ads**

» Falsified performance measurement or outcomes such as viewability measurement, app installs, and location

» **Sophisticated bots**

image source: Integral Ad Science

General Invalid Traffic (GIVT) is a term used to describe traffic that is considered invalid or fraudulent, but can be detected and filtered out through standard fraud detection techniques. GIVT includes benevolent non-human traffic from sources such as bots, spiders, and other automated programs that are working to index content from all over the internet. The purpose of these bots is to understand web pages better, so that the information can be more easily retrieved when needed.

As mentioned, GIVT can be filtered out using some standard techniques (such as IP blacklisting, user agent filtering, and device fingerprinting) but obviously can still cause advertisers to waste money on non-human traffic. It is generally considered less harmful than more sophisticated forms of ad fraud, which can be more difficult to detect and prevent.

Sophisticated Invalid Traffic (SIVT) is a term used to describe a more advanced form of ad fraud that is designed to evade traditional fraud detection and prevention techniques. SIVT typically involves more sophisticated techniques and tactics, such as domain spoofing, ad stacking, or click injection, that are specifically designed to generate fraudulent ad impressions or clicks while appearing to be legitimate user behaviour.

Unlike General Invalid Traffic (GIVT), which can be detected and filtered out using standard fraud detection techniques, SIVT requires more advanced and sophisticated fraud detection methods. For example, SIVT may involve the use of advanced bot networks that are designed to mimic human behaviour, or the use of complex algorithms that are able to evade standard fraud detection models.

SIVT can be particularly damaging to advertisers, as it can result in a significant waste of ad spend and can negatively impact the performance of ad campaigns. It’s important for advertisers, publishers, and ad tech companies to implement advanced fraud detection and prevention measures to protect their ad campaigns from SIVT, and to stay up-to-date on the latest fraud tactics and techniques in order to stay ahead of criminals.

different forms of ad fraud

3

There is a very wide range of different types of digital advertising fraud, and we have included a list of these below with a short explainer for each.

Impression fraud

Impression fraud occurs when an ad is served but never actually seen by a human. This is often done by loading ads in hidden iframes or bots generating fake impressions.

Click fraud

Click fraud occurs when an individual or bot repeatedly clicks on an ad to generate revenue or exhaust an advertiser's budget.

Invisible pixels

Invisible pixels are small, hidden pixels that are used to track user behaviour without being noticed by users. Fraudsters can use invisible pixels to generate fake impressions or clicks without users realising they are being tracked.

Invisible ads

Invisible ads are ads that are displayed to users but are invisible to the naked eye, either because they are hidden behind other content or because they are displayed in a colour that matches the background of the webpage. Invisible ads can generate fraudulent impressions and clicks without being noticed by users.

Ad stacking

Ad stacking involves stacking multiple ads on top of each other so that only the top ad is visible to the user. This allows publishers to generate more revenue from a single loading.

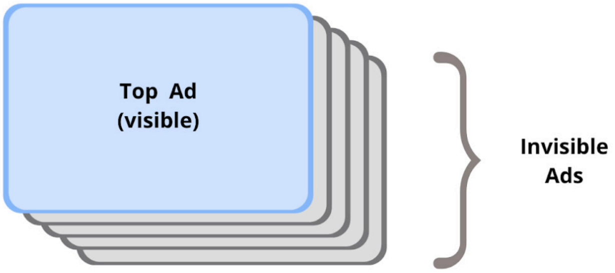

image source: Edgemesh

Cookie stuffing

Cookie stuffing involves adding cookies to a user's browser without their consent, which can artificially inflate ad metrics such as click-through rates and/or conversions.

Identity stuffing

This is the practice of substituting, overwriting or manipulating publisher identifiers during the RTB processes to reduce the effectiveness of user consent signals in an attempt to improve performance.

Domain spoofing

Domain spoofing is a tactic used by fraudsters to make it appear as though an ad impression is coming from a legitimate website, when in fact it is coming from a completely different source. This can be accomplished through a variety of techniques, such as using fake URLs, redirect chains, or using domain names that are similar to legitimate websites.

Fake leads

Advertisers may also receive fake leads or sign-ups generated by fraudsters using fake or stolen identities.

Ad injection

Ad injection involves the unauthorised injection of ads into a user's browser or mobile device, often through browser extensions or malware.

Click injection

Click injection is a type of mobile ad fraud that involves fraudsters injecting clicks into an app at the moment that it is installed on a user's device. This can artificially inflate click-through rates and deceive advertisers into paying for fraudulent clicks.

Ad hijacking

A well-known issue particularly in paid search - which is also known as brand poaching, direct linking or URL jacking. This is when an affiliate impersonates a brand by running ads that look identical to a brand's ads. Consumers typically have no idea they have clicked on an affiliate ad, and the affiliate is paid a commission for conversions that the brand would have captured directly.

Pixel stuffing

Pixel stuffing occurs when a publisher places an ad in a small pixel size on a webpage, making it difficult or impossible for users to see. This method allows them to artificially inflate ad impressions.

Geo-targeting fraud

Geo-targeting or geo-masking involves falsifying a user's location to deceive advertisers into paying for targeted ads delivered to low-quality traffic disguised as higher quality traffic from a more sought-after country or region.

Click Flooding

Fraudsters send large numbers of fraudulent click reports in the hopes of delivering the last-click prior to installs to get paid out by advertisers. This can result in advertisers making ill-informed decisions based on false insights, invalid installs and misattribution of ad spend.insights, invalid installs and misattribution of ad spend.

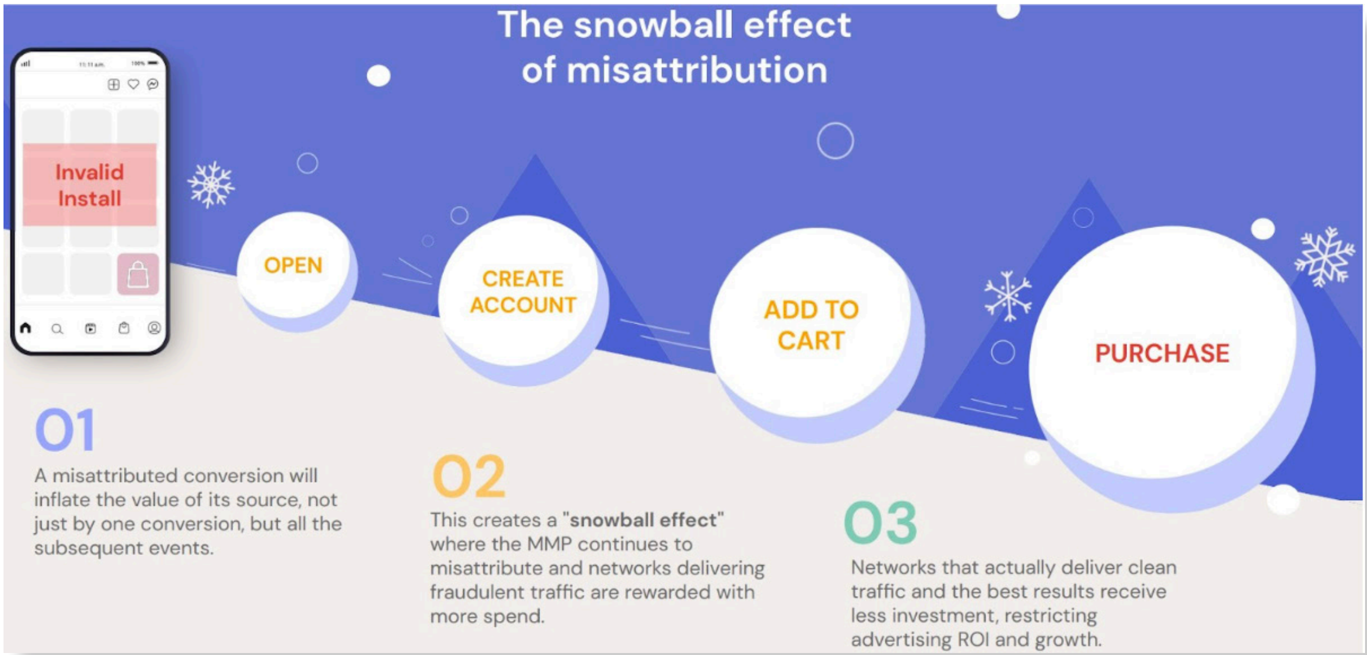

The snowball effect of misattribution

01

A misattributed conversion will inflate the value of its source, not just by one conversion, but all the subsequent events.

02

This creates a **"snowball effect"** where the MMP continues to misattribute and networks delivering fraudulent traffic are rewarded with more spend.

03

Networks that actually deliver clean traffic and the best results receive less investment, restricting advertising ROI and growth.

image source: TrafficGuard

App install fraud

App install fraud occurs when fraudsters use bots or incentivised downloads to drive app installs. This can be used to artificially inflate app download metrics and deceive advertisers into paying for ads that are not generating genuine user engagement.

App install farms

App install farms are real locations with real people using real devices to install apps. It's a basic manual way to generate the activity they're paid for, but in regions where human resources are cheaper it can still be profitable. Scammers will change the IP address and refresh device IDs for each outcome so that the process can be relentlessly repeated without flagging any warnings if not monitored well enough.

Click farms

As per the above, click farms are groups of individuals who are paid to click on ads, either manually or using automated scripts. Click farms can generate a large volume of fake clicks and can be difficult to detect and prevent, particularly if the clicks are the result of real human 'clickers' who demonstrate genuine human behaviours. Again, the scammers will regularly change the IP address and refresh device IDs.

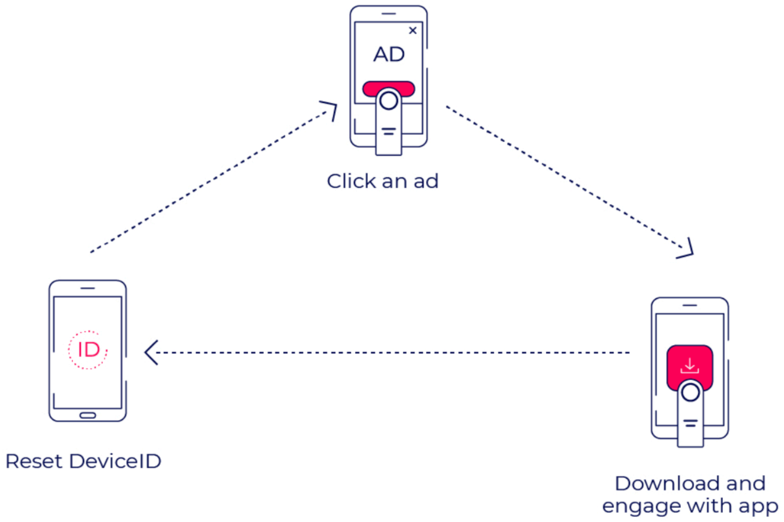

image source: AppsFlyer

SDK Spoofing

SDK spoofing is a type of mobile ad fraud where criminals hack software development kits (SDKs) generating fake installs and in-app events on real devices in order to manipulate advertising by mimicking user behaviour.

Ad laundering

Ad laundering involves funnelling ad inventory through a network of intermediaries to obscure its origin and deceive advertisers into paying for low-quality or fraudulent placements.

Viewability fraud

Viewability fraud occurs when an ad is served but never actually seen by a human. This can be done by placing ads below the fold or behind other content on a webpage, or by loading ads in small, hidden iframes.

Click spamming

Click spamming is a technique used by fraudsters to generate a large number of clicks on an ad in a short period of time, overwhelming the ad server and preventing legitimate clicks from being recorded. This can be done using botnets or other techniques that allow the fraudsters to generate a high volume of clicks in a short amount of time.

Redirect attacks

This is where fraudsters redirect users that have clicked on one ad to multiple other ads and back again in the blink of an eye. All of those redirects count as individual clicks, generating revenue for the cybercriminals.

Cross-device fraud

Cross-device fraud occurs when fraudsters use multiple devices to generate fake clicks or impressions on ads. This can be done by using botnets or other techniques to mimic human behaviours across multiple devices, making it difficult to detect and prevent.

Bot traffic

This refers to non-human traffic generated by bots, crawlers, or javascripts that mimic human behaviours. This can be used to generate fake clicks, impressions, or conversions.

Botnet traffic

Botnets are networks of compromised computers or devices that are controlled by a single entity, often for the purpose of generating fake ad impressions or clicks. Botnets can be difficult to detect because they use a large number of IP addresses and user agents to mimic human behaviour. $15m USD was seized from Swiss bank accounts belonging to operators behind the infamous '3ve' online ad fraud scheme, which had seen 1.7m devices infected with the Kovter botnet between 2015-2018.

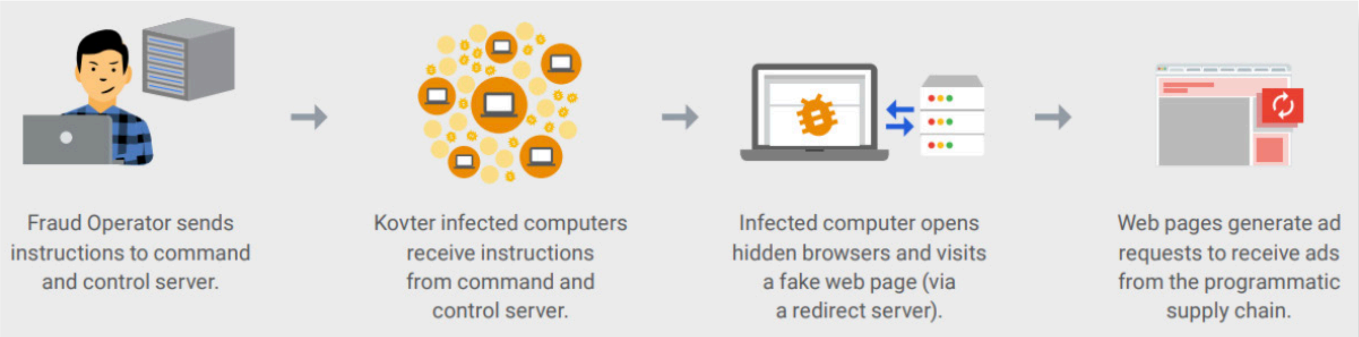

Fraud Operator sends instructions to command and control server.

Kovter infected computers receive instructions from command and control server.

Infected computer opens hidden browsers and visits a fake web page (via a redirect server).

Web pages generate ad requests to receive ads from the programmatic supply chain.

image source: Bleeping Computer

Fraudulent data center traffic

Data center traffic is any traffic whose IP address shows it's originating from a server rather than common devices such as PCs and smartphones. A fraudulent data center is essentially a network of servers, each having its own IP address which well organised fraudsters can use to send out massive volumes of bots to direct the advertiser's campaign traffic and run up fake impressions. Then, to avoid detection, the bot will switch to a new IP address for each visit to the campaign. Detecting fake ad traffic from data centers can be tricky because the traffic is sometimes indistinguishable from real traffic.

Device fingerprinting

Device fingerprinting involves tracking and identifying users based on unique characteristics of their devices, such as their browser settings, operating system, and installed plugins. Fraudsters can use this information to generate fake clicks or impressions that appear to come from legitimate devices.

Domain doppelgangers

Domain Doppelganging is a technique used by fraudsters to register a domain name that is very similar to a legitimate one. For example it may be missing a dot between host/subdomain and domain. The fraudsters then host fake websites on the fake domain and generate fraudulent traffic or clicks on ads that appear on the site.

Programmatic ad fraud

Programmatic fraud occurs when fraudsters take advantage of the real-time bidding process used in programmatic advertising to generate fake impressions or clicks on ads. This can be done by manipulating bid requests or injecting fraudulent ad tags into the bidding process.

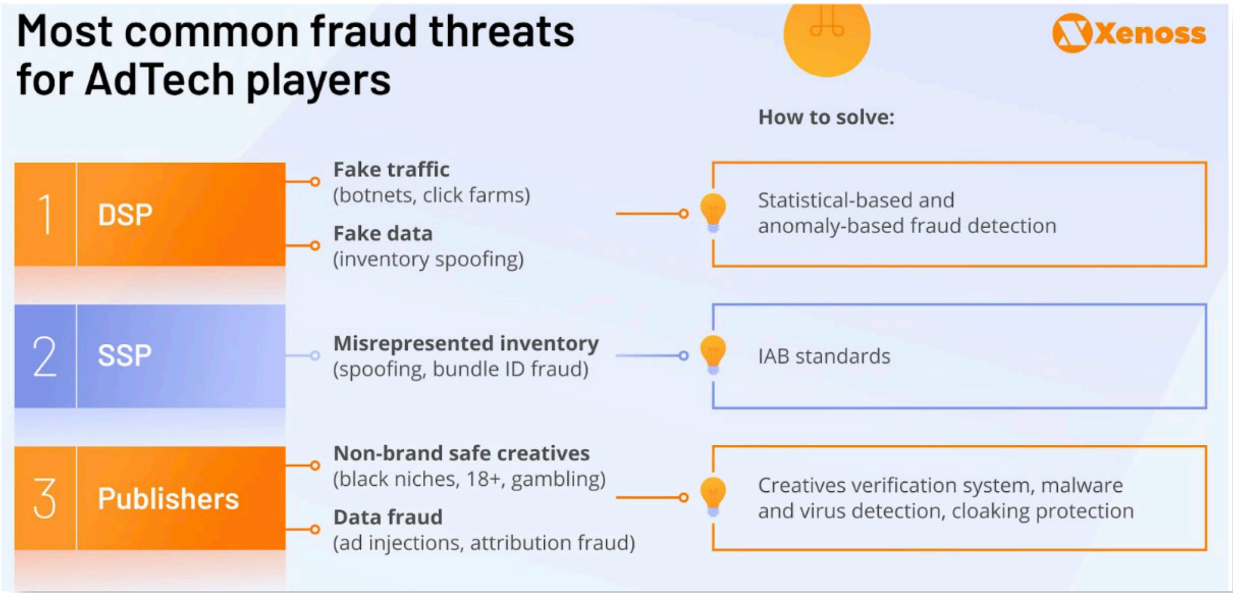

Most common fraud threats for AdTech players

		How to solve:
1 DSP	**Fake traffic** (botnets, click farms) **Fake data** (inventory spoofing)	Statistical-based and anomaly-based fraud detection
2 SSP	**Misrepresented inventory** (spoofing, bundle ID fraud)	IAB standards
3 Publishers	**Non-brand safe creatives** (black niches, 18+, gambling) **Data fraud** (ad injections, attribution fraud)	Creatives verification system, malware and virus detection, cloaking protection

image source: Xenoss

Spoofing ads.txt files

This involves fraudsters spoofing the domain of a legitimate publisher and adding it to their own ads.txt file (more on ads.txt and app-ads.txt further on in this document). This makes it appear as though the fraudulent inventory is authorised by the legitimate publisher, making it more difficult for advertisers to detect and prevent.

Cookie Syncing

Cookie syncing involves matching the cookies of different ad tech platforms to track user behaviour across multiple domains. This can be used by fraudsters to generate fake clicks or impressions that appear to come from legitimate sources.

Domain Chaining

Domain chaining is a technique used by fraudsters to chain together a series of fake websites to generate fake clicks or impressions. This can be done by using redirects or other methods to move users from one fake website to another, making it difficult to track the source of the fraudulent traffic.

Misrepresented traffic

Sellers may misrepresent the source of their traffic in order to sell more ad inventory or to charge higher prices for their ad space. As a result advertisers may be paying for traffic that is not as valuable as they believe it to be.

Audio SSAI falsification

As uncovered by DoubleVerify (more info here) Audio is not immune to issues with fraud. SSAI falsification works similarly to SSAI fraud on CTV. Fraudsters begin by spoofing residential IP addresses and audio apps, whilst setting up fake SSAI servers to falsify audio ad requests and pass these into SSPs at usually very cheap prices in order to attract demand from buyers.

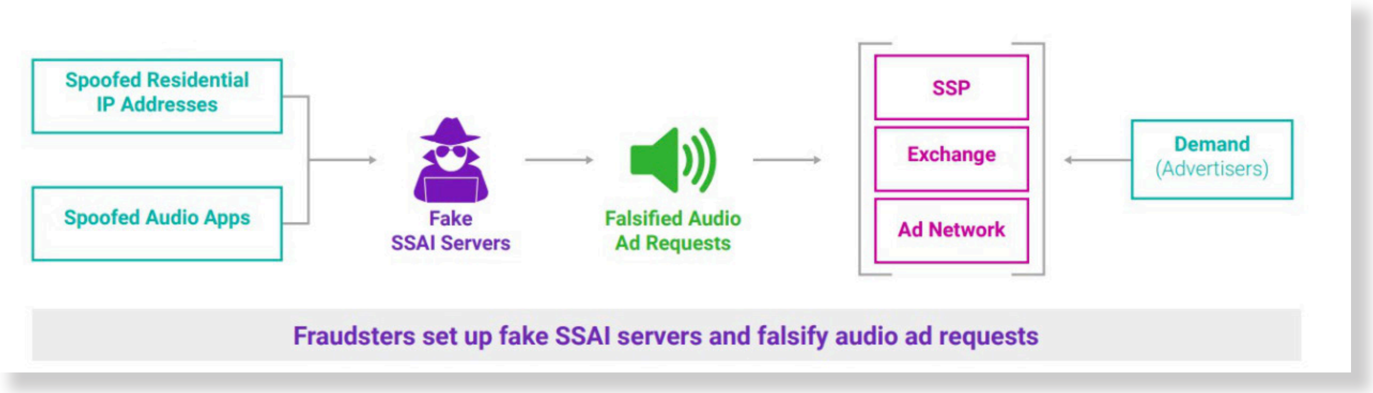
Spoofed Residential IP Addresses
Spoofed Audio Apps
Fake SSAI Servers
Falsified Audio Ad Requests
SSP
Exchange
Ad Network
Demand (Advertisers)
Fraudsters set up fake SSAI servers and falsify audio ad requests

image source: DoubleVerify

recommendations & best practices 4

We have separated this section into several parts to include recommendations for sellers & buyers, as well as how to competently use analytics to audit activity and a section on IAB Tech Lab's transparency standards.

Advertisers can implement a range of different methods to ensure their investments are protected as best as possible against "bad actors" within the supply chain. This can be done via a number of different mechanisms:

DSP Features

Platforms like Yahoo DSP have actively built proprietary fraud protection schemes meeting MRC spec requirements. These schemes include some of the below approaches and also include anti-fraud guarantees to provide funding back to advertisers against fraudulent traffic.

Standard exclusion & inclusion lists

These can be implemented at a seat, advertiser or campaign level. Creating standard exclusion lists (or blocklists) helps provide a top layer protection from fraud. An inclusion list is a list of sites or domains and bundle IDs that an advertiser considers to be safe, acceptable, and trustworthy environments to serve ads to. Several factors are considered when adding sites or domains to an inclusion list, including the brand safety and brand suitability requirements of the product or service being advertised.

Real-time scoring

Creating a set of risk rules to determine the likelihood of fraud helps protect your advertisers in real time.

Client-side forensics

Custom javascript is used to help determine true-domain, browser spoofing, unattended browser instances, and other signals that can lead to further detection and elimination of fraudulent traffic.

Supply transparency and supply slicing

Technology leverages core browser APIs to detect true-domains and thereby determine domain spoofing. This information is then married with incoming bid requests to determine slices of inventory that should be blocklists to prevent further instances of spoofed inventory.

Manual traffic review

A team of traffic analysts who monitor traffic signals to determine and block suspicious traffic patterns and behaviours.

Supply policy enforcement

All domains and page content are subjected to automated classifiers to detect content that violates supply partner policies. Based on the classifier confidence these domains are subjected to further manual review and banned if required.

Post-bid blocking

Advertisers can run post-bid protection capabilities from media quality vendors such as Integral Ad Science, DoubleVerify and Oracle in both programmatic and direct campaigns. Post-bid blocking is primarily deterministic, occurring after an impression is won. It detects IVT in real-time, with blocking decisions based on what's happening in the moment and is a very effective approach for saving on costs and managing brand safety.

Pre-bid targeting

Pre-bid targeting excludes segments that have been identified as high IVT domains or unsafe ad environments and is primarily designed to maximise programmatic ad budget efficiency. The segments are usually set up within a DSP enabling advertisers to set parameters around campaigns or line items and bid only on the most relevant incoming requests. There is a built-in level of uncertainty in pre-bid segmentation as this targeting is applied prior to a bidding decision and is inherently probabilistic.

SPO

Advertisers and agencies should have greater control over their supply strategies through Supply Path Optimisation. This can be achieved through DSP proprietary tools, through commercial agreements with SSPs, as well as through activation against only ads.txt or ads.cert enabled partners. Also agencies should insist upon always using separate trading seats for each of their clients so as to improve reporting and minimise bid duplications.

Use a reputable ad network or platform

Choose a reputable ad network or platform that has a good track record in preventing Ad Fraud.

Use Ad Fraud detection tools

Utilise Ad Fraud detection tools, such as HUMAN (formerly White Ops), Forensiq, Fraudlogix and The Media Trust , to identify suspicious traffic patterns, click fraud, and other fraudulent activities.

Monitor your campaigns

Keep a close eye on your ad campaigns and regularly monitor your traffic sources. If you notice any unusual or suspicious activity, take immediate action.

Set up fraud filters

Set up filters to block fraudulent traffic, such as bots and click farms, from accessing your ads. This can help reduce the risk of Ad Fraud.

Use viewability metrics

Utilise viewability metrics to measure how many users actually see your ads. This can help identify potential fraudsters who may be generating fake impressions.

Work with trusted partners

Partner with trusted vendors and publishers who have a history of providing high-quality traffic and transparent reporting.

Educate yourself

Stay informed about the latest Ad Fraud trends and best practices. Attend industry conferences and events, read industry publications, and connect with other advertisers and marketers to share insights and learn from each other.

Industry benchmarks and discrepancies

These can vary widely depending on the type of ad campaign, the industry, and the region. Fraud can artificially inflate metrics such as impressions and clicks, leading to misleading results and wasted ad spend. Therefore, it is crucial to implement the above measures to prevent Ad Fraud and protect your advertising budget.

Bot Protection for Large Publishers

There are enterprise solutions for large publishers that are designed to help website, app and application owners detect bot traffic that are attempting to access their sites and/or services. These solutions will also often provide clean data on traffic patterns, types, and volume to train machine learning algorithms to be more accurate in detecting bots and non-human traffic of all forms.

These types of services generally use multiple patented technologies to detect and mitigate bots where they make initial contact rather than allowing them to reach websites or apps first. Many also leverage the capabilities of the network effect as these platforms work interconnectedly to adapt to any re-tooling attempts across thousands of different domains and apps.

Be aware that these types of solutions will incur a range of different costs, and some may need to be overseen and managed by competent engineering staff.

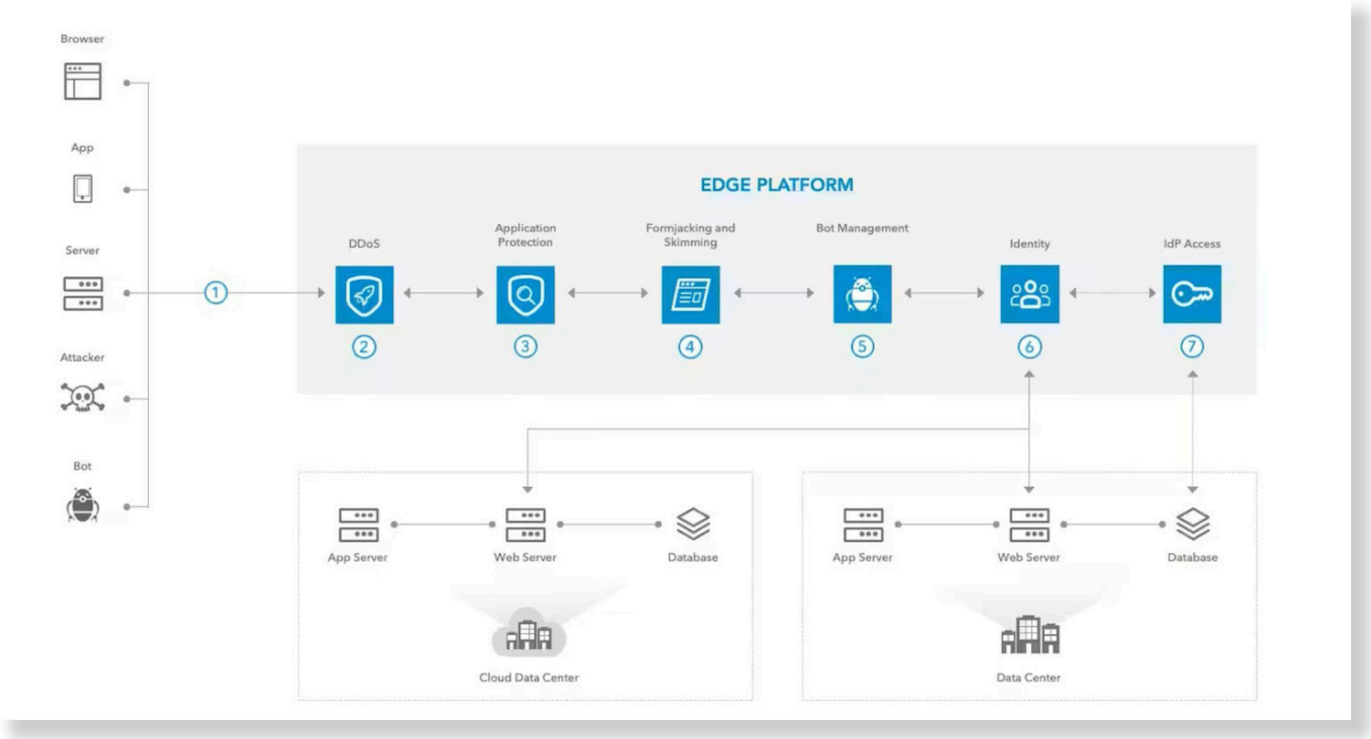

image source: Akamai Technologies

Using Analytics

For this section we have included content from Dr. Augustine Fou, the owner of FouAnalytics - as he is a leading expert in leveraging analytics to interrogate digital marketing campaigns and run digital media audits to mitigate ad fraud. The intention of including this selected content from a guest contributor is not to promote any one specific solution, but rather to recommend the overall analytical approach and promote a willingness to be prepared to dive deeper into the critical details.

The content below is a very small selected sample of his approach and has been sourced with permission from the two articles below:

Forensic Auditing of Digital Media:

https://www.linkedin.com/pulse/forensic-auditing-digital-media-augustine-fou/

How Scrutinise Clicks from Programmatic Campaigns:

https://www.linkedin.com/pulse/how-use-fouanalytics-scrutinize-clicks-from-programmatic-fou/

Interrogating Delivery Reports

Without any specialized tools or technologies, advertisers should ask their media agencies for DSP reports, by month, or preferably by day. This tells you how many bids you won and paid for. Then ask for ad server reports, by month, or preferably by day. This tells you how many ads were served. Normally "ads served" should be one-to-one with "bids won;" when you win the bid you have the right to serve your ad into the ad slot.

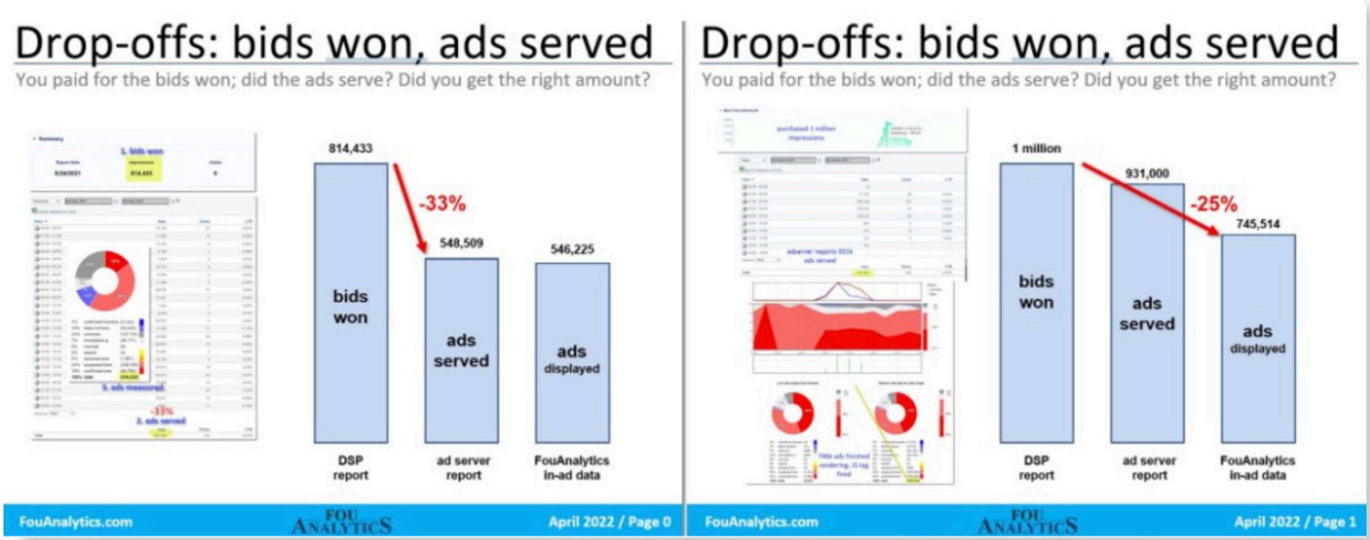

Drop-offs: bids won, ads served

You paid for the bids won; did the ads serve? Did you get the right amount?

FouAnalytics.com — April 2022 / Page 0

Drop-offs: bids won, ads served

You paid for the bids won; did the ads serve? Did you get the right amount?

FouAnalytics.com — April 2022 / Page 1

image source: FouAnalytics

The two slides above show scenarios that are so typical it's hard to believe. The slide on the left shows 814k bids won, but only 549k ads served, a 33% "drop off." That means a third of the ads you paid for did not even serve. The slide on the right shows that you got 93% ads served (931k compared to 1 million), but only 745k of the ads rendered on screen and were measured by FouAnalytics, a 25% "drop off." So you got 1/4 less than you paid for.

Why does this "drop off" happen? Bot activity and fraud. Bad guys make their bots more efficient by skipping the loading of the ad. If they already get paid for the bids won, why waste time and bandwidth loading the ad. Bots just move on to the next fake bid.

In addition to the simple stuff above, by adding a FouAnalytics tag into the ads, we can check to see if the ads were rendered on screen. Ads can be served (sent out) from the ad server, but never arrive in the device to be displayed on screen. This problem is especially bad in mobile because lower bandwidth means the ads may not arrive in time, before the user leaves the page or scrolls away.

Video ads are far larger in file size than display ads, so the likelihood of video ads never rendering on screen is even higher. So it's always good to check how many ads were actually displayed on screen. The FouAnalytics tag is javascript; and it is set to fire "async" (asynchronously). That means it will execute when the ad finishes rendering. This is how we can tell the ad was displayed. By comparing the counts measured by a FouAnalytics tag with the numbers from the ad server, you can see the second "drop off" in the drop-off slides above.
As an advertiser, the only ads that are valuable to you are the ones that are served and also displayed on screen. Are you getting what you paid for?

Interrogating Click Patterns

Click pattern charts in FouAnalytics make it easy to tell which clicks are humans versus bots, just by looking at the scatterplot of the x,y coordinates of the clicks. Clicks are grouped by screen resolutions. The top 3 in the slide below are desktop monitors and laptop screens in landscape orientation. The bottom 2 on the left are mobile screen resolutions in portrait orientation. Note the bottom right, with the large red circles, larger circles means multiple clicks on the same x,y coordinates. Bots can easily click the same location over and over; but it is hard to get a whole bunch of humans to click on the exact same x,y coordinates as other humans. See the screenshots further below from on-site measurement. It is visually clear which clicks are from bots and which clicks are from humans.

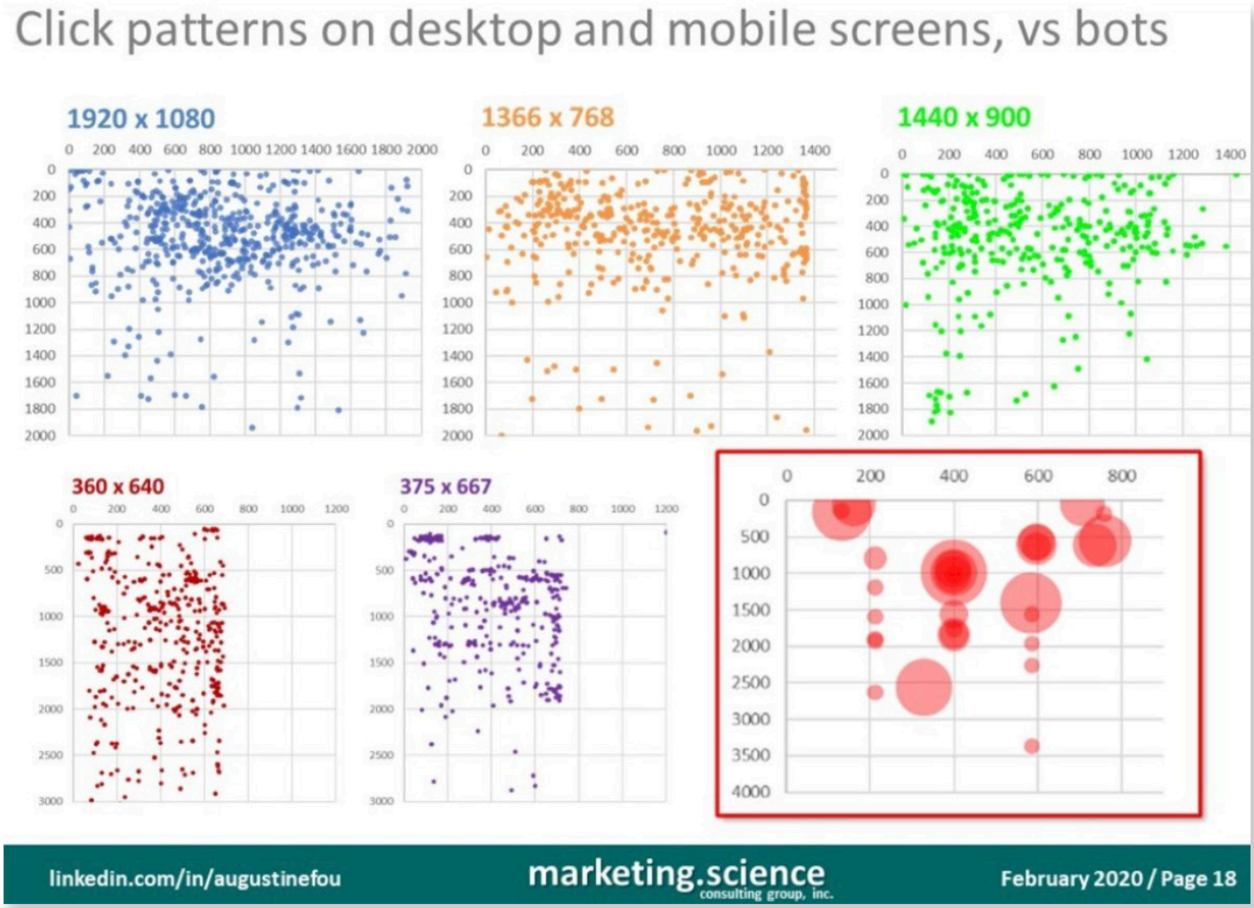

Click patterns on desktop and mobile screens, vs bots

linkedin.com/in/augustinefou

marketing.science consulting group, inc.

February 2020 / Page 18

image source: FouAnalytics

The key takeaway for advertisers is that having lots of clicks is meaningless if a full 85 - 99% of the clicks are bots. Even though you got lots of clicks, your campaigns did not actually perform well, because bots don't convert. Clicks are easy to report, and therefore are usually the metric that gets reported to advertisers. Advertisers also think that "more is better" (larger numbers of clicks are better) and often simply assume their campaigns are performing better because there are more clicks and higher click rates.

Google Ads - Confirmed Click Feature

Confirmed Click is a Google Ads product feature, first introduced in April 2021, that adds a confirmation to ad placements that may be generating accidental clicks. When a user clicks an ad with Confirmed Click, the user will be asked to confirm their intent to visit the advertiser page by clicking a button associated with an action, such as 'Visit site'.

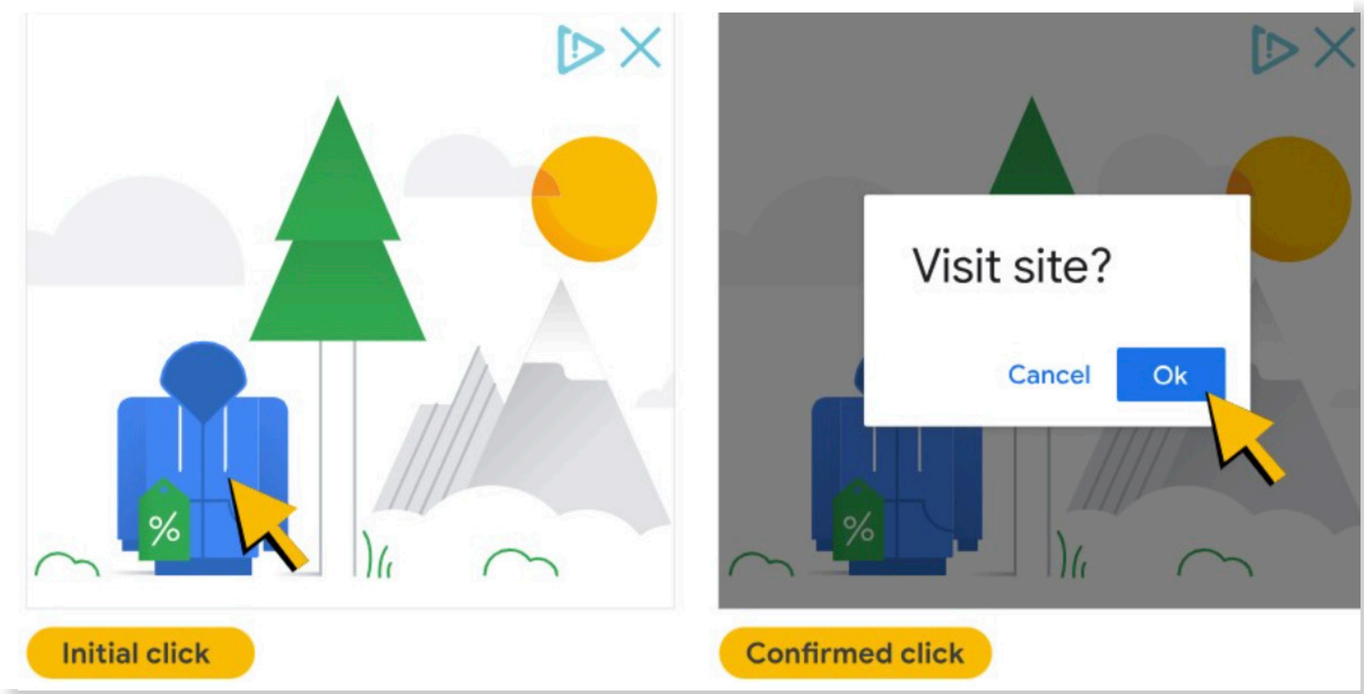

image source: Google

Confirmed Click is applied when a Google Ads system determines that the ads on the impacted sites are generating unintended clicks. This creates a poor user experience because the user ends up on advertiser landing pages instead of their intended content.

By introducing the second click, Google improved the experience by letting the user confirm their intent to visit the advertised page. Typically, users that intended to click on the ad will still click through to view the ad. To find out more about this feature click here

IAB Tech Lab's Programmatic Transparency Standards

IAB Tech Lab has developed several standards that are aimed at improving transparency and trust in the programmatic advertising ecosystem. These have not been developed as a 'silver bullet' solution for fraud but do help build trust through improved transparency in the supply chain.

There are three key standards that publishers should be aware of:

ads.txt

The IAB Tech Lab's ads.txt standard is a simple text file that is placed on a publisher's website. It lists the authorised sellers of the publisher's inventory, which helps to prevent counterfeit inventory and domain spoofing. This can help to ensure that advertisers are buying inventory from legitimate sources and not from counterfeiters.

The app-ads.txt provides the same information for publishers to share within their in-app environments.

ads.txt and app-ads.txt support transparent programmatic digital media transactions and can remove the financial incentive from selling counterfeit and misrepresented media.

Similar to robots.txt, the ads.txt file can only be posted to a domain by a publisher's webmaster, making it valid and authentic. As a text file it is easy to update and maintain, making it very flexible. The data required to populate the file is readily available in the OpenRTB protocol, making it simple to gather and target.

Because publishers sell their inventory through a variety of sales channels, ads.txt supports the following types of supplier relationships:

- **Domain owners who sell on exchanges through their own accounts**
- **Networks and sales houses who programmatically sell on behalf of domain owners**
- **Content syndication partnerships where multiple authorised sellers represent the same inventory**

ads.txt works by creating a publicly accessible record of authorised digital sellers for publisher inventory that programmatic buyers can index and reference if they wish to purchase inventory from authorised sellers.

First, participating publishers must post their list of authorised sellers to their domain. Programmatic buyers can then crawl the web for publisher ads.txt files to create a list of authorised sellers for each participating publisher. Then programmatic buyers can create a filter to match their ads.txt list against the data provided in the OpenRTB bid request.

Example:

Example.com publishes ads.txt on their web server listing three exchanges as authorised to sell their inventory, including Example.com's seller account IDs within each of those exchanges.

http://example.com/ads.txt:

#< SSP/Exchange Domain >, < SellerAccountID >, < PaymentsType >, < TAGID >
greenadexchange.com, 12345, DIRECT, AEC242
blueadexchange.com, 4536, DIRECT
silverssp.com, 9675, RESELLER

Note: The seller's Publisher.ID will be specified in the 'SellerAccountID' field in the ads.txt.

A buyer receiving a bid request claiming to be example.com can verify if the exchange and SellerAccountID matches the authorised sellers listed in example.com/ads.txt file.

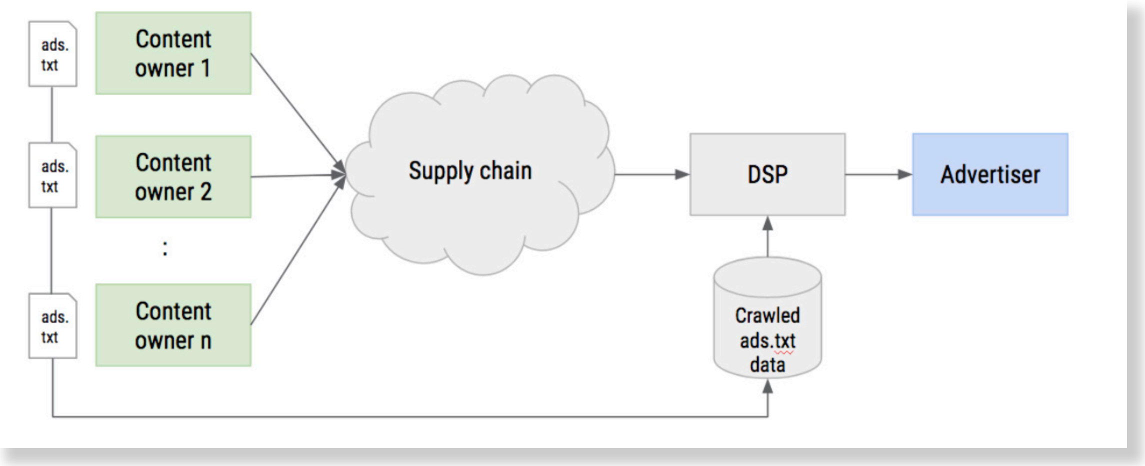

Source: https://iabtechlab.com/ads-txt-about

Creating, hosting and maintaining an accurate ads.txt file is an important responsibility for a publisher ad ops team. For a good working example see https://www.news.com.au/ads.txt

Sellers.json

Sellers.json is an extension of the ads.txt standard and it provides more details about the authorised sellers that are listed in the ads.txt file. SSPs and exchanges post their sellers.json files for buyers to review and the format is readable by both humans and machines. The file includes information such as the seller's name, website, and contact information. It also includes the relationship between the seller and the publisher, such as whether the seller is a direct or reseller.

This is not a file that publishers will have to maintain but it's important to be cognisant of the process that SSPs and exchanges go through and to know how to review the relevant .json files to ensure that they are correct and that the domains they are authorised to sell are being properly represented therein.

For a good working example see https://pubmatic.com/sellers.json

SupplyChain Object (SCO)

The SupplyChain Object (SCO) is a standard that provides a way for buyers to report and validate the supply chain of an impression, including information such as the publisher, the seller, and any intermediaries involved in the transaction. This standard aims to increase transparency around the supply chain and allow for better tracking of inventory and more efficient supply chain management.

Again this is not a process that publishers need to manage directly, but it's good to be aware of and a reminder of how important it is to maintain accurate and up-to-date ads.txt and app-ads.txt files.

All of these standards were created by IAB Tech Lab to improve transparency and trust in the programmatic ecosystem. By providing more detailed information about the authorised sellers and supply chain of an impression, these standards can help to reduce fraud and improve the efficiency of programmatic advertising.

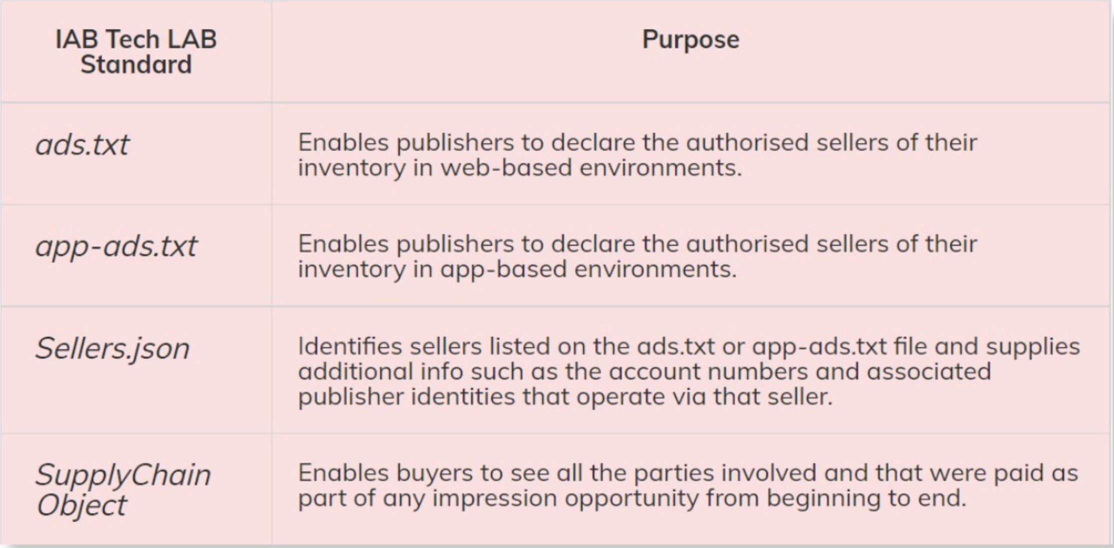

IAB Tech LAB Standard	Purpose
ads.txt	Enables publishers to declare the authorised sellers of their inventory in web-based environments.
app-ads.txt	Enables publishers to declare the authorised sellers of their inventory in app-based environments.
Sellers.json	Identifies sellers listed on the ads.txt or app-ads.txt file and supplies additional info such as the account numbers and associated publisher identities that operate via that seller.
SupplyChain Object	Enables buyers to see all the parties involved and that were paid as part of any impression opportunity from beginning to end.

image source: https://iabaustralia.com.au/resource/open-rtb-update-iab-tech-lab-sept-2019/

IAB Tech Lab's Best Practices for the Disclosure of Ad Fraud Attacks

In May 2022 IAB Tech Lab published a set of proposed Ad Fraud Disclosures to try and improve how ad fraud attacks are disclosed. The intention was to consolidate a set of best practices that can drive alignment and consistency across the industry when participants choose to disclose information about an ad fraud attack.

Using these best practices will help the ecosystem by providing clear, useful and timely information that allows the ecosystem to combat ad fraud in an effective manner. Despite the challenges that ad fraud represents across various dimensions, collaboration and responsible disclosures will enable the industry to further secure and harden the ecosystem against attacks.

Traditionally ad fraud attacks are disclosed in bespoke and sometimes fragmented ways, leading to both a lack of consistency and clarity in what is disclosed, as well as potential gaps in both understanding and verification of attacks. By implementing the best practices included in this proposal, disclosing entities will help the industry to better address and respond to ad fraud attacks in a mutually beneficial way.

For more information on these best practices, please access this IAB Tech Lab document here

industry case studies 5

We have 6 case studies in total

Click below to access the full report

DoubleVerify Exposes ViperBot

DoubleVerify

Video Ad Fraud Executed via malicious Creatives

The Media Trust

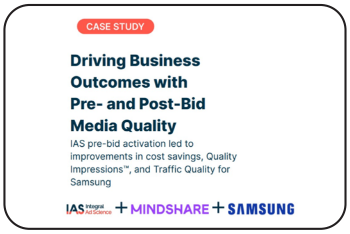

CASE STUDY

Driving Business Outcomes with Pre- and Post-Bid Media Quality

IAS pre-bid activation led to improvements in cost savings, Quality Impressions™, and Traffic Quality for Samsung

IAS Integral Ad Science + MINDSHARE + SAMSUNG

IAS Pre-Bid Activation for Samsung

Integral Ad Science

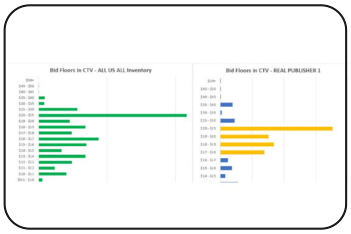

CTV Bid Floors Analysis

FouAnalytics

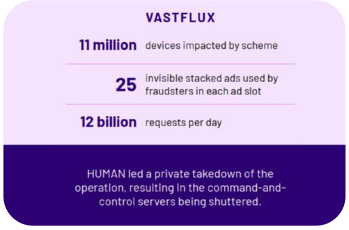

VASTFLUX

11 million	devices impacted by scheme
25	invisible stacked ads used by fraudsters in each ad slot
12 billion	requests per day

HUMAN led a private takedown of the operation, resulting in the command-and-control servers being shuttered.

The VASTFLUX ad fraud operation

HUMAN

Optimising Google AdSense by Fighting Ad Fraud for an Online Education Provider

Optimising Google AdSense to Fight Ad Fraud

TrafficGuard

There are two main drivers for ad fraud in CTV:

The first is the high-value transactions that CTV CPMs yield. According to the IAB Tech Lab, The CTV market represents tens of billions of USD in ad spend and has CPMs x25 greater than web ads, making it a very attractive channel for fraudsters and bad actors to target.

The second driver is that CTV environments are signal poor which means that the signals that these devices omit across the supply chain make it very difficult for buyers and ad verification vendors to identify whether the signal being transmitted is by a real person, device, or a bot.

The reason for this poor signalling is that the current IVT detection techniques available were designed for the web and don't work as well in CTV - creating the perfect conditions for fraudsters and their schemes.

Lastly, the video supply chain is convoluted and involves many different players. In order for buyers to transact safely in the supply pool, it requires both sides to demand higher standards in securing the Over the top (OTT) supply chain.

Different Types of CTV Ad Fraud

There are many different types of ad fraud; one of the most common categories are fake requests known as "spoofing" which is the act of disguising a signal from a fraudulent activity as being from a legitimate source and is the most common type of CTV fraud.

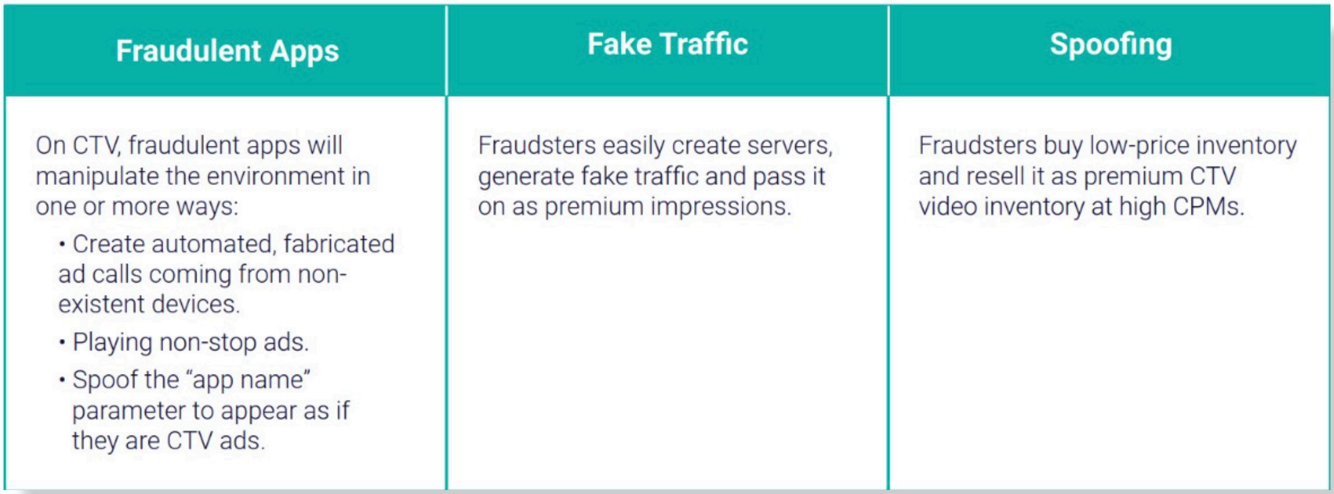

Fraudulent Apps	Fake Traffic	Spoofing
On CTV, fraudulent apps will manipulate the environment in one or more ways: • Create automated, fabricated ad calls coming from non-existent devices. • Playing non-stop ads. • Spoof the "app name" parameter to appear as if they are CTV ads.	Fraudsters easily create servers, generate fake traffic and pass it on as premium impressions.	Fraudsters buy low-price inventory and resell it as premium CTV video inventory at high CPMs.

image source: DoubleVerify

Bad actors disguised as audiences who are watching ads on CTV devices and take away ad spending through spoofing tactics, such as:

IP spoofing

Fraudsters disguise themselves as specific computer IP addresses to impersonate another computer system and create fraudulent traffic.

Device spoofing

Fraudsters disguise the actual device they're using as different devices, such as mobile phones, browsers, or laptops, to click on ads and fill out forms on websites. Usually, a large number of clicks coming from a single device would indicate fraud. But by disguising their actual device, bad players can pose as audiences and take away marketers' ad spend without getting notified of fraud.

SDK spoofing

A type of mobile fraud where bad actors impersonate genuine app installs, purchases, and clicks.

Server-side-ad-insertion (SSAI) spoofing

SSAI spoofing happens when fraudsters send out millions of ad requests from data centers pretending to be CTV devices. SSAI is a technology developed by publishers whereby ads are "stitched" into OTT videos so customers no longer have to wait for ad players to launch. As SSAI is relatively new, fraudsters look for loopholes in the technology to launch fraud operations.

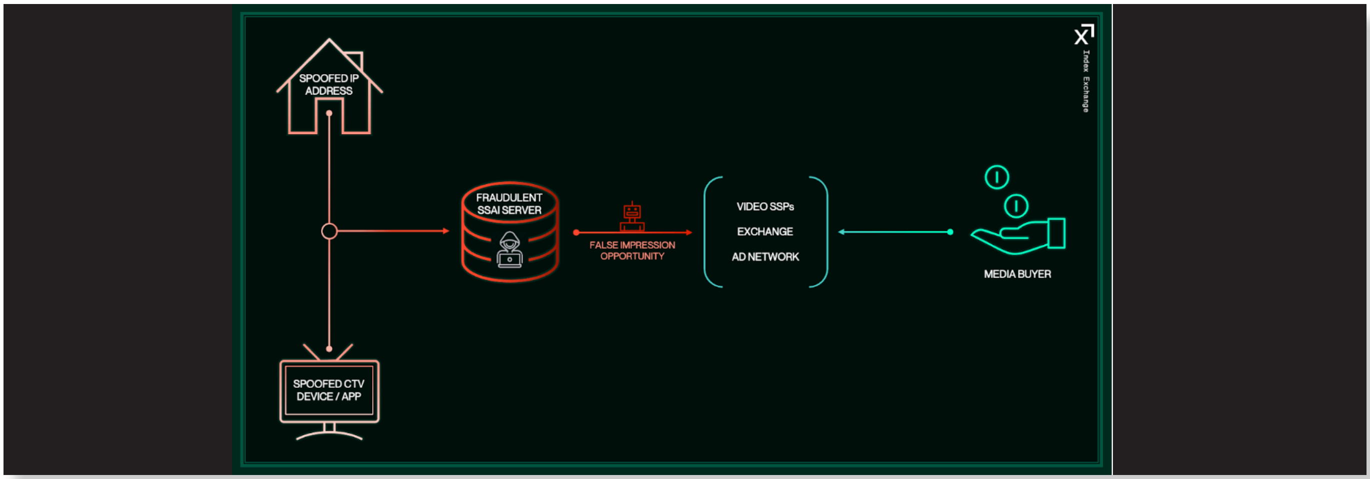

image source: DoubleVerify

Another common type of ad fraud is the misinterpretation of inventory where bad actors buy low-value inventory and disguise it as high value. There have been many of these scandals, however the most recent cross-device scandal called DiCaprio uncovered how fraudsters tricked advertisers into thinking they were buying video advertisements on Roku-connected TV devices.

Ultimately, bad actors work to disguise themselves as legitimate users, devices, and signals to trick advertisers into buying illegitimate inventory across devices.

What are the Prevention Tactics?

Existing tactics and protocols

There are a number of simple checks and prevention tactics that buyers and sellers can utilise including:

Simple Checks:

IP Address

Ensuring that the same IP Address is used across requests will help to weed out any fraudulent activity.

Click timestamp and action timestamp

The click timestamp is the time when someone arrives on a site after clicking an ad. The action timestamp is the time when that person completed an action on a site. In most cases, if an IP address has a number of click timestamps but no action timestamps, then that is likely click fraud.

User-agent

Useful for identifying whether someone on a particular IP is the same person. It takes note of all the features of the device being used to access a website or app such as type of computer or device, internet browser, software, and more.

IAB Tech Lab Standards

IAB Tech Lab's Authorised Digital Sellers (ads.txt and app-ads.txt) standards were covered previously in this document and are important as they give buyers a list of authorised sellers of legitimate inventory. These specifications are a simple way to help ensure that the media inventory being bought and sold, especially across applications, is legitimate.

Currently the majority of Australian CTV buys remain pre-packaged products and audiences delivered directly to buyers (depending upon their requirements) by premium partners and vendors without having to expose brands to any of the risks of open-market programmatic bidding environments. It's simple, highly brand-safe, competently productised, and very well-tailored to marketers' requirements.

However, as programmatic CTV demand continues to grow – we have to expect that these CTV/OTT products will become progressively democratised as buyers inevitably seek to access this supply as they see fit. The industry is obligated to ensure that we can safely enable these capabilities across the full range of a marketer's appetite for risk. Hence the ads.cert specifications.

Ads.cert2.0

These are available today and include a set of protocols with cryptographic signatures that are intended to be used to authenticate the server-side ad integration (SSAI), the device, the user and the app. An example of this is a protocol called authenticated connections for SSAI bid requests which identifies that the party sending the ad request is who they claim to be. This relies on the owner of the domain to push out a cryptographic key that allows buyers to identify the source of the inventory is legitimate.

These standards are initially focused on SSAI (Server-Side Ad Insertion) transactions specifically as recent security research has highlighted schemes where parties have attempted to impersonate SSAI platforms. These schemes are challenging to identify, as traffic appears to originate from the same cloud platforms and hosting providers that service genuine SSAI businesses.

Due to the multi-pronged set of requirements, there are several protocols wrapped into this, so we'll break them down one-by-one for you:

Call Signs protocol

Allows a company to accurately identify other companies involved in a specific ad transaction, thanks to Domain Name System records.

Authenticated Connections protocol

Gives both advertisers and publishers confidence in the authenticity of the origin of any requests, thereby preventing interference in server-to-server requests.

Authenticated Delivery protocol

Authenticates the data in a given bid request, allowing buyers and sellers to see if the price or location of a given bid has been tampered with.

Authenticated Devices protocol

Attests to the legitimacy of the device on which a given ad is being served. An example of this is the recent watermark developed by Roku which ensures that all impressions come from the device instead of the SSAI which enables all notifications to come with a specialised watermark proving the legitimacy of the inventory being sold.

IAB Tech Lab has specifically recommended that all SSAI providers immediately implement the Authenticated Connections protocol in particular – and advises both buyers and sellers to start insisting upon this protocol for any CTV transactions as soon as it becomes available to them.

For the full set of specifications for ads.cert 2.0 from IAB Tech Lab simply click here

Fraud Prevention Vendors

Enlisting the help of companies like HUMAN and Integral Ad Science work to safeguard media owners and media buyers from digital attacks, including bots, fraud, and account abuse. Ultimately, they help to protect a media owner's revenue and ensure media buyers are purchasing quality inventory with their media budgets.

Future Developments

The IAB Tech lab will continue to work on adding additional protocols to the existing standards mentioned above however the biggest area for development is around measurement.

The Open Measurement Software Development Kit (OM SDK) is designed to facilitate third-party viewability and verification measurement for ads served to web video and native app environments and is now available in CTV.

As the specification develops, the OM SDK functionality with the support of device manufacturers can help to detect whether a connected TV is on/off which can try to prevent ad fraud.

Additionally, in late 2022, the IAB Tech Lab unveiled the Advanced TV roadmap outlining what developments will be focused on in the next several years. For example, the release of on-the-glass measurement and attribution gives buyers real-time data and the confidence to know that a real person viewed the ad on a particular device.

conclusion 7

The intention of this handbook is to highlight the importance of this important topic and ensure that all industry participants take it seriously and look to 'get their hands dirty' with the related details to ascertain what requirements, partnerships and internal processes may be needed to mitigate as many of the associated risks as possible.

In terms of benchmarking there are various MRC-certified sources that monitor the percentage of IVT (note that not all will separate the insights into GIVT and SIVT) along with other key verification metrics. Some are regional, some cover specific markets and some are more detailed than others - and it's advisable to keep across all the key trends as much as possible.

Examples of the two most common sources of regular insights for the Australian market are below and we recommend regular reviews of this data to compare with whatever you are seeing at your end.

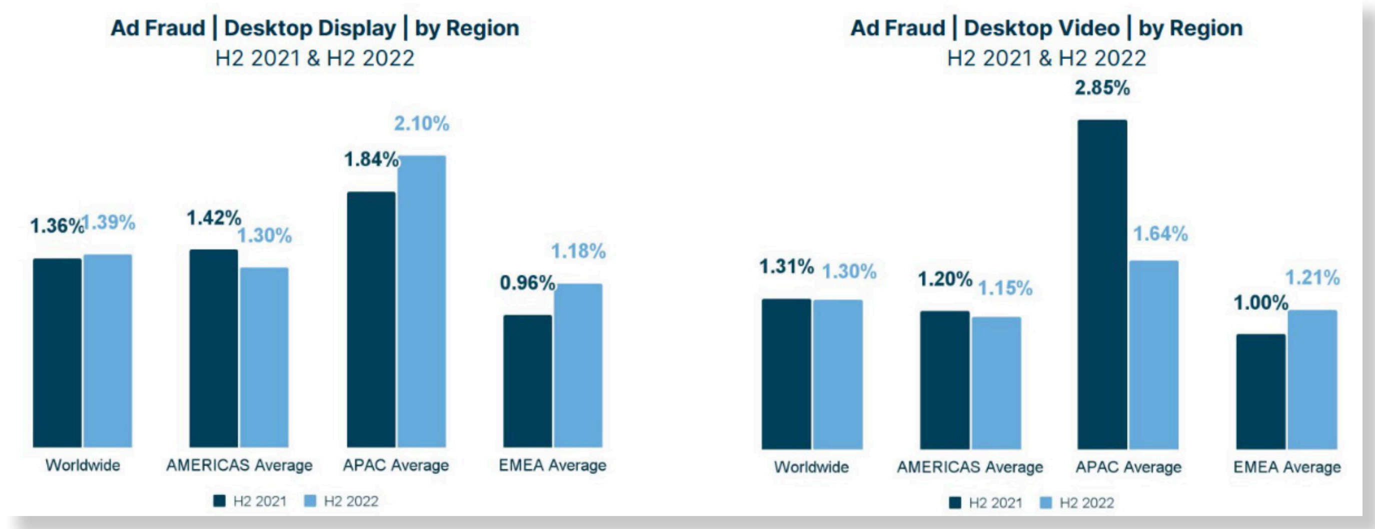

image source: Integral Ad Science (Media Quality Report)

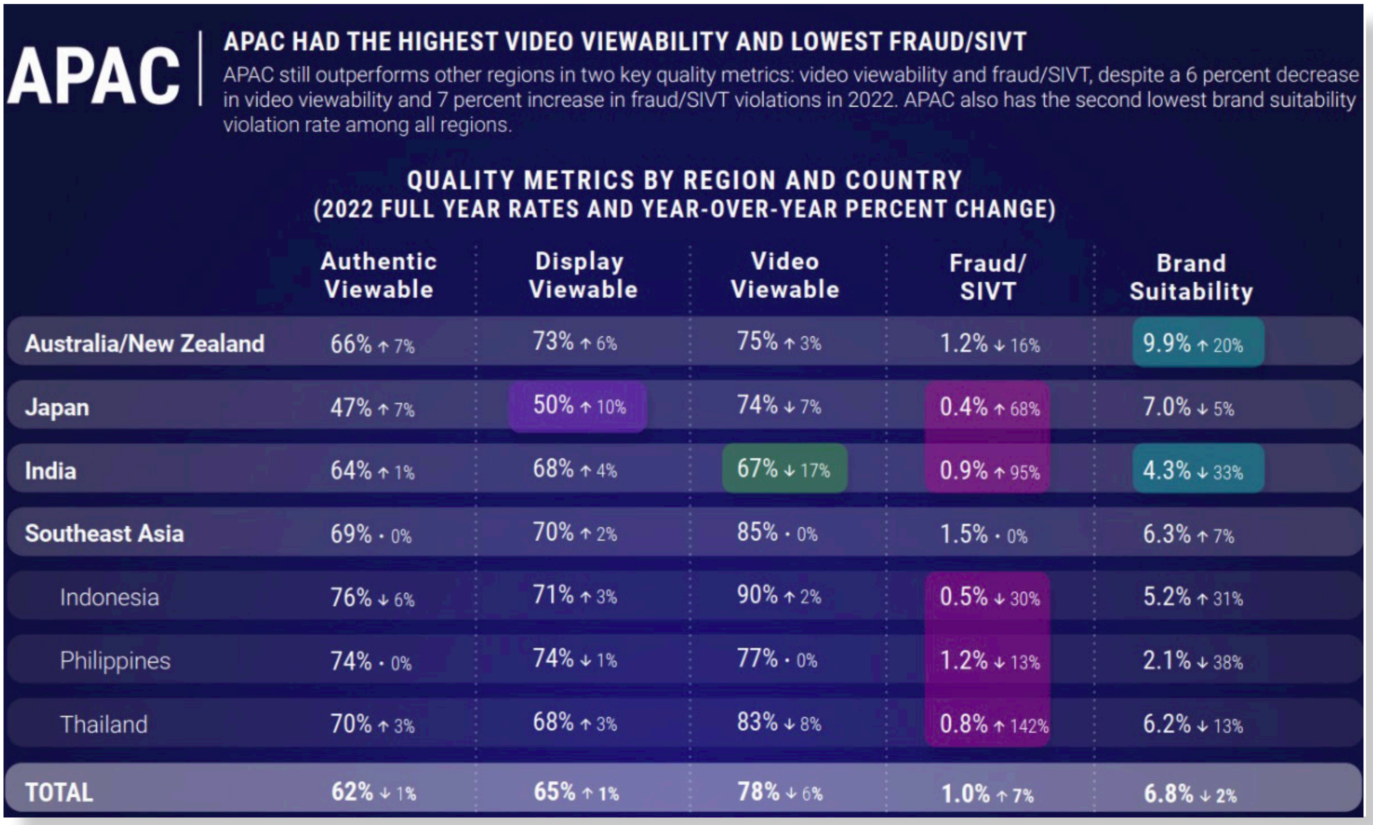

APAC

APAC HAD THE HIGHEST VIDEO VIEWABILITY AND LOWEST FRAUD/SIVT

APAC still outperforms other regions in two key quality metrics: video viewability and fraud/SIVT, despite a 6 percent decrease in video viewability and 7 percent increase in fraud/SIVT violations in 2022. APAC also has the second lowest brand suitability violation rate among all regions.

QUALITY METRICS BY REGION AND COUNTRY
(2022 FULL YEAR RATES AND YEAR-OVER-YEAR PERCENT CHANGE)

	Authentic Viewable	Display Viewable	Video Viewable	Fraud/ SIVT	Brand Suitability
Australia/New Zealand	66% ↑ 7%	73% ↑ 6%	75% ↑ 3%	1.2% ↓ 16%	9.9% ↑ 20%
Japan	47% ↑ 7%	50% ↑ 10%	74% ↓ 7%	0.4% ↑ 68%	7.0% ↓ 5%
India	64% ↑ 1%	68% ↑ 4%	67% ↓ 17%	0.9% ↑ 95%	4.3% ↓ 33%
Southeast Asia	69% · 0%	70% ↑ 2%	85% · 0%	1.5% · 0%	6.3% ↑ 7%
Indonesia	76% ↓ 6%	71% ↑ 3%	90% ↑ 2%	0.5% ↓ 30%	5.2% ↑ 31%
Philippines	74% · 0%	74% ↓ 1%	77% · 0%	1.2% ↓ 13%	2.1% ↓ 38%
Thailand	70% ↑ 3%	68% ↑ 3%	83% ↓ 8%	0.8% ↑ 142%	6.2% ↓ 13%
TOTAL	62% ↓ 1%	65% ↑ 1%	78% ↓ 6%	1.0% ↑ 7%	6.8% ↓ 2%

image source: DoubleVerify (Global Insights Report)

A current hot topic is AI, and it's probably worth touching upon this as a final note. Looking forward, it's inevitable that these capabilities will be promoted by many as a killer solution for fighting fraudsters. Unfortunately , it is more likely that it will only add to the problems we are facing.

Sophisticated fraud is essentially a capabilities arms race and it often feels like vendors are constantly playing catch-up. Machine Learning is not new (however it's packaged) and can only work with the data inputs it receives, which will largely be the methods of known schemes which have already been uncovered. AI is simply parasitical in that sense, but the hope is that by integrating both supervised and unsupervised Machine Learning models, vendors can gradually improve their capabilities.

A supervised model is trained on an extensive set of properly labelled data inputs. During fraud detection, each transaction is classified as either fraud or non-fraud. Over time, the Machine Learning model creates an algorithm that distinguishes fraudulent transactions. The problem is that we live in a dynamic world where everything changes very quickly. Fraudsters are very good at inventing new fraud techniques and In these cases, supervised models can be ineffective because they are constantly being exposed to phenomenons they are unfamiliar with.

Unsupervised models however, are trained using unlabeled data. Therefore, unsupervised models should be better at detecting newer forms of fraud schemes. These models can detect behavioural anomalies by identifying transactions that do not conform to the majority.

However, my honest fear is that the capabilities of AI will be more effectively leveraged by sophisticated criminals for fraudulent purposes. The key reasons for my concern here are:

Speed and efficiency

AI can process large amounts of data and perform tasks quickly, which makes it a potentially useful tool for automating fraudulent activities.

Anonymity

AI can be used to carry out fraudulent activities without leaving a traceable human trail.

Evasion

AI can be used to evade detection by constantly generating fake or misleading information that is difficult for humans to detect as fraudulent.

Generating fake or misleading content

AI could be used to create fake websites, social media accounts, or other online content at scale designed to deceive. This could include generating fake reviews or manipulating online ratings to mislead consumers.

Automation of scams

AI could be used to automate scams or fraudulent schemes at scale.

Increased sophistication

AI could be used to increase the sophistication of fraudulent schemes and cyber-attacks by adapting more quickly to the defences as they evolve.

Improved impersonation

The use of AI to impersonate real people is becoming an increasingly sophisticated and effective form of deception.

So, whilst we have to be hopeful that these capabilities will help to improve things as technologies advance we also have to be cognisant of the problem only getting worse, and hence this is an area for us to try and keep an eye on as things evolve.

We hope that you have found this guidance useful and if you have any feedback at all, or further case studies, please feel free to email us directly at jonas@iabaustralia.com.au

further reading 8

Integral Ad Science's Media Quality Reports
https://integralads.com/insider/category/resources/mqr/

DoubleVerify's 2023 Global Insights Report
https://doubleverify.com/2023-global-insights-report/

Google: Invalid Activity Guidance
https://www.google.com/ads/adtrafficquality/invalid-activity/

Google: Protecting your ad supported CTV experiences at scale
https://goo.gle/protecting_ctv_experiences

Google: 2022 Ads Safety Report
https://services.google.com/fh/files/misc/2022_google_ads_safety_report.pdf

Index Exchange: Exchange Quality Guidance
https://www.indexexchange.com/product/exchange-quality/

IAB Tech Lab: Ad Fraud Disclosures
https://iabtechlab.com/ad-fraud-disclosures/

IAB Tech Lab: ads.cert 2.0 Authentication Protocols
https://iabtechlab.com/ads-cert/

HUMAN: Enterprise Bot Fraud Benchmark Report
https://www.humansecurity.com/2023-enterprise-bot-fraud-benchmark-report

TrafficGuard: Five Signs of Ad Fraud in your Digital Marketing Campaigns
https://www.trafficguard.ai/resources/five-signs-of-ad-fraud-in-your-digital-marketing-campaigns