Ad Fraud Essentials

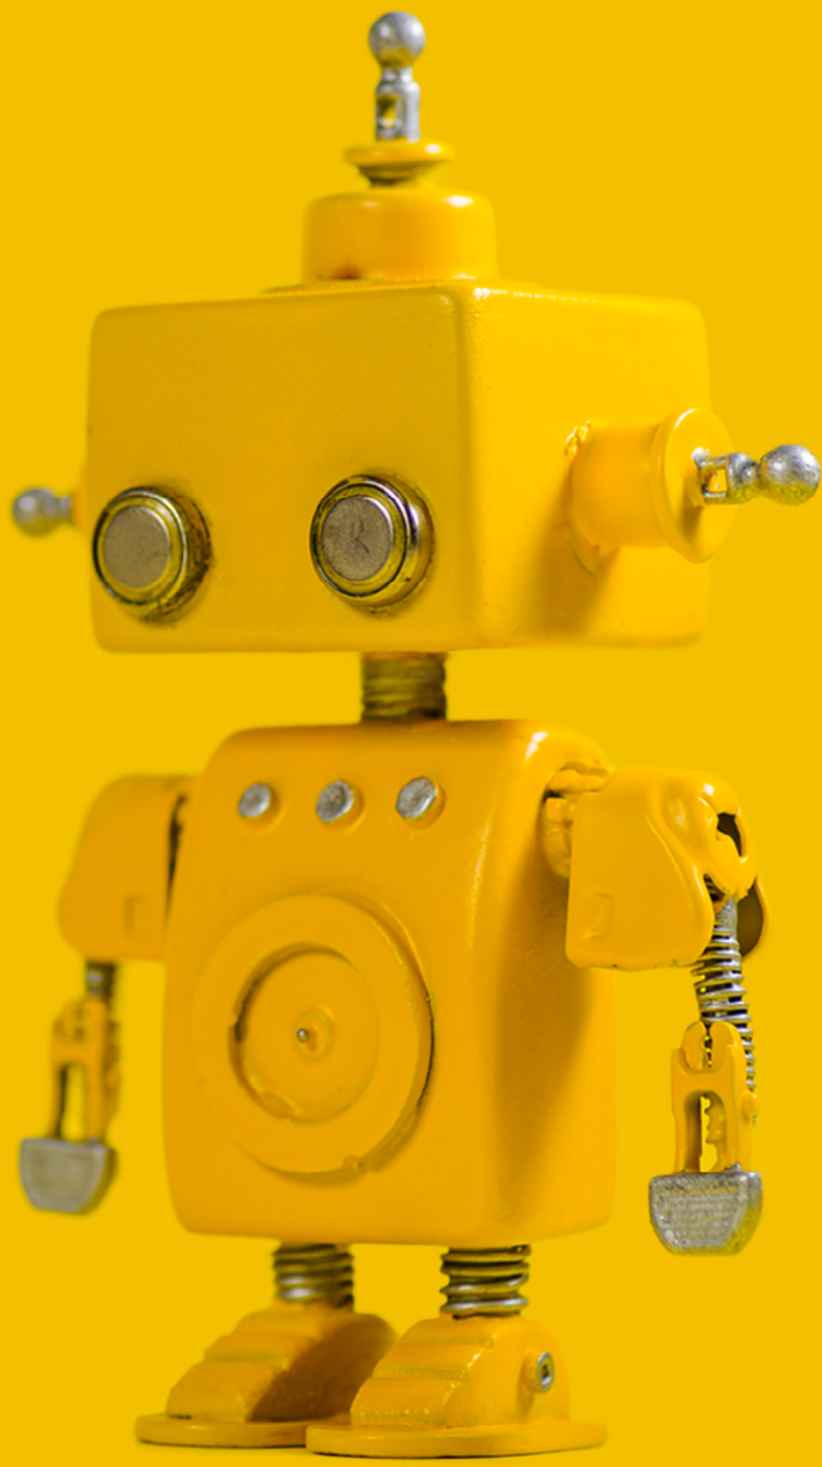

IAS Integral Ad Science

Table of contents

INTRODUCTION

According to the World Federation of Advertisers (WFA), it is estimated that by 2025, over $50 billion will be wasted annually on ad fraud.

As advertising spend continues to shift to digital media, ad fraud has proven to appear in many different forms: bots, pixel stuffing, malicious apps, and more. Ad fraud has evolved in such a way that it can impact every element of digital advertising. Even though it is pervasive, there continues to be general confusion around how detrimental ad fraud is for the digital advertising ecosystem.

It is apparent that ad fraud has grown into a worldwide multibillion dollar business. At the same time, digital has become the medium of choice for many consumers and advertisers. With digital ad spending predicted to reach nearly $500 billion globally by 2025, the industry is more focused than ever before on combating digital ad fraud.

In our latest Industry Pulse report, IAS surveyed 915 members of the digital advertising industry — agencies, brands, publishers, and ad tech partners — to better understand their concerns for the year ahead. Since 2017, ad fraud has been top of the priority list for brands to address. In the latest report, 42.8% of all respondents agreed that ad fraud will continue to be a main priority.

Top 3 industry priorities in 2019 (% of total respondents)

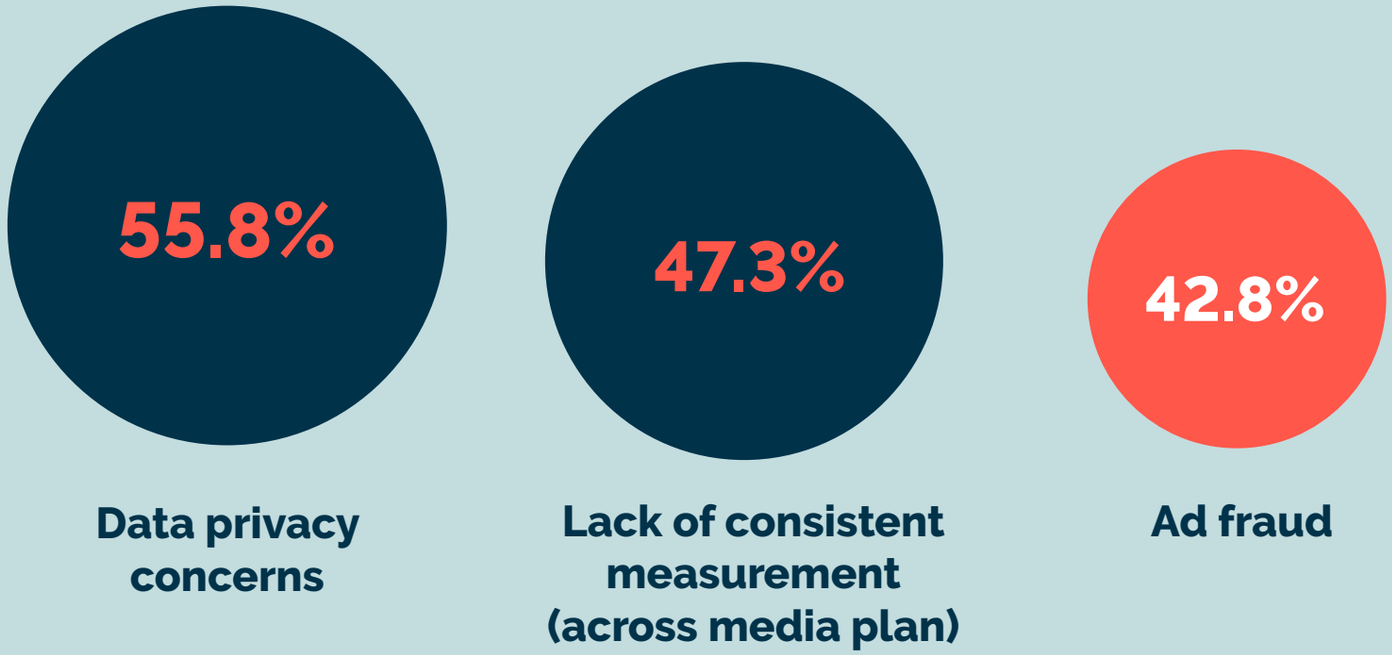

CHAPTER ONE

Overview

IAS defines digital **ad fraud** as any deliberate activity that prevents the proper delivery of ads to the intended audience, in the intended place. Ad fraud directly results in financial loss or opportunity loss by the advertiser and/or publisher in a particular transaction.

Ad fraud is often referred to as **invalid traffic (IVT)**. Invalid traffic is a broad term describing online activity that does not always come from a real user, therefore the impressions do not represent legitimate advertising consumption.

The industry has categorized invalid traffic in two ways: General Invalid Traffic (GIVT) and Sophisticated Invalid Traffic (SIVT).

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**

The bottom line is, any IVT that generates actions that take away from the proper delivery of an ad can impact every type of advertising — display, video, mobile, audio, search, and social. The landscape of fraud is always evolving and can impact all players in the ecosystem — even the most premium publishers. Every traffic source requires constant re-evaluation.

So, what *exactly* is ad fraud?

AD FRAUD IMPACTING ADVERTISERS

- Delivering pre-roll video placements in display banner slots
- Hiding ads behind or inside other page elements so they can't be viewed by a user
- Serving ads on a site other than the one requested in an **RTB** request
- Falsifying user characteristics like location and browser type
- A network of malicious **bots** that infiltrate user browsers and devices to generate impressions under false identities

AD FRAUD IMPACTING PUBLISHERS

- Malicious bots or fraudulent apps that sell inventory automatically
- Frequently refreshing the ad unit or page to hinder a user's opportunity to engage with the ad

What *isn't* ad fraud?

- Web crawlers and other good bots
- Collisions (ads from the same brand accidentally appearing on the same page)
- Muted video
- Poor viewability

How pervasive is ad fraud?

Fraud fluctuates depending on the country and the device. Each day, IAS measures over one trillion media quality metrics globally. Leveraging this dataset, our **Media Quality Report** provides both the buy and sell sides of the industry with benchmarks and insight into the levels of ad fraud in major digital markets.

From these reports, IAS sees that fraud fluctuates from 0.4% - 11.7% depending on whether buyers implement an ad fraud prevention or detection strategy. Globally, non-optimized campaigns that do not utilize any ad fraud mitigation attract the most fraud.

Global ad fraud levels for desktop display

WORLDWIDE NON-OPTIMIZED
14.7 % **11.7 %**

WORLDWIDE OPTIMIZED
0.8 % **1.1 %**

H1 2018 **H1 2019**

OPTIMIZED AGAINST AD FRAUD

AUSTRALIA
0% 5%
0.9% **1.5%**

CANADA
0% 5%
0.8% **0.8%**

FRANCE
0% 5%
0.4% **0.9%**

GERMANY
0% 5%
0.7% **1.2%**

ITALY
0% 5%
0.6% **0.8%**

OPTIMIZED AGAINST AD FRAUD

JAPAN
0% 5%
1.6% **2.6%**

NEW ZEALAND
0% 5%
0.4% **0.4%**

SPAIN
0% 5%
0.5% **0.6%**

UK
0% 5%
0.7% **0.9%**

USA
0% 5%
1.8% **1.9%**

% of total impressions

Ad fraud rates are on the rise

Ad fraud has become, and will continue to be, a major issue for the advertising industry. From H2 2018 to 2019, IAS's Media Quality Reports show that ad fraud levels have increased. With content creation accelerating at an unprecedented rate and advertisers investing more money in digital, fraudsters will follow the money and capitalize on lucrative opportunities.

For more highlights on the state of ad fraud in digital advertising around the world, check out our latest Media Quality Reports.

Global non-optimized fraud levels

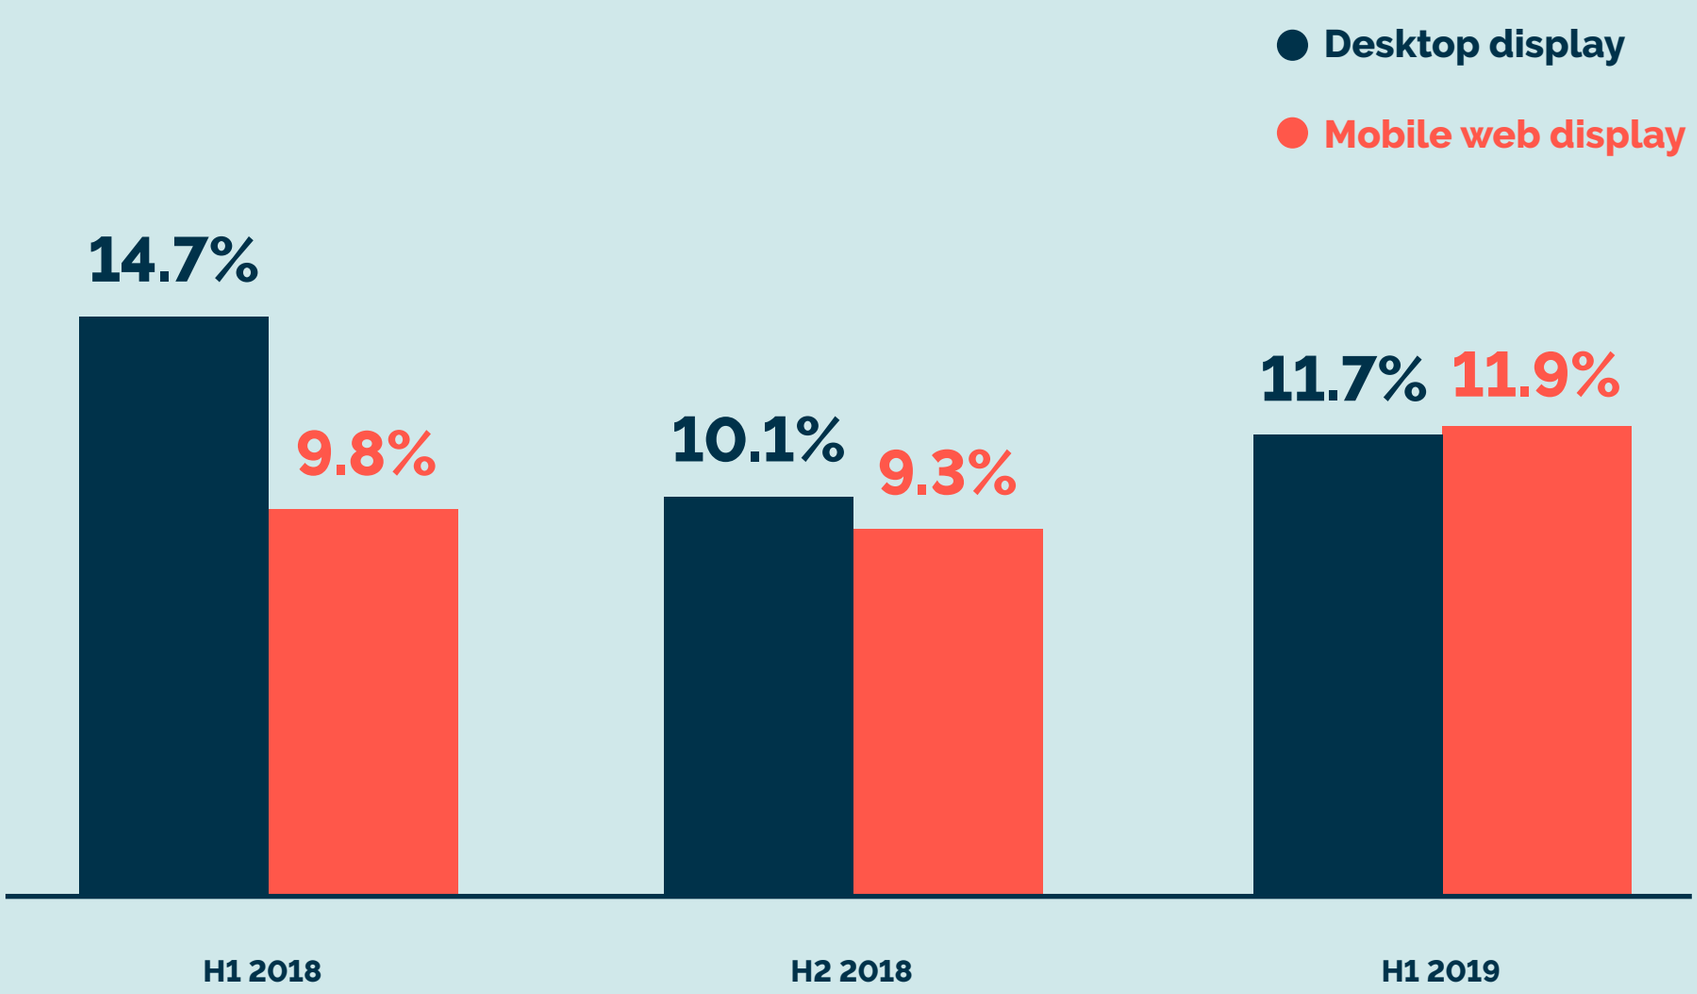

Source: Integral Ad Science, Media Quality Reports H1 2017 - H1 2019

Why does ad fraud occur?

Ad fraud is a readily available and straight-forward operation for those with the right skills and knowledge. To carry out criminal activity across the web, all a fraudster needs is a few computers and the technical know-how. This low barrier to entry makes ad fraud an attractive prospect over other forms of fraud that would require more investment or infrastructure.

Within the realm of criminal activity, ad fraud has incredibly high payout potential at a low risk to the fraudster, as it is difficult to penalize by law. Most ad fraud is committed in countries with indifferent or ineffective cybercrime law enforcement.

Payout potential vs. risk in the fraud ecosystem

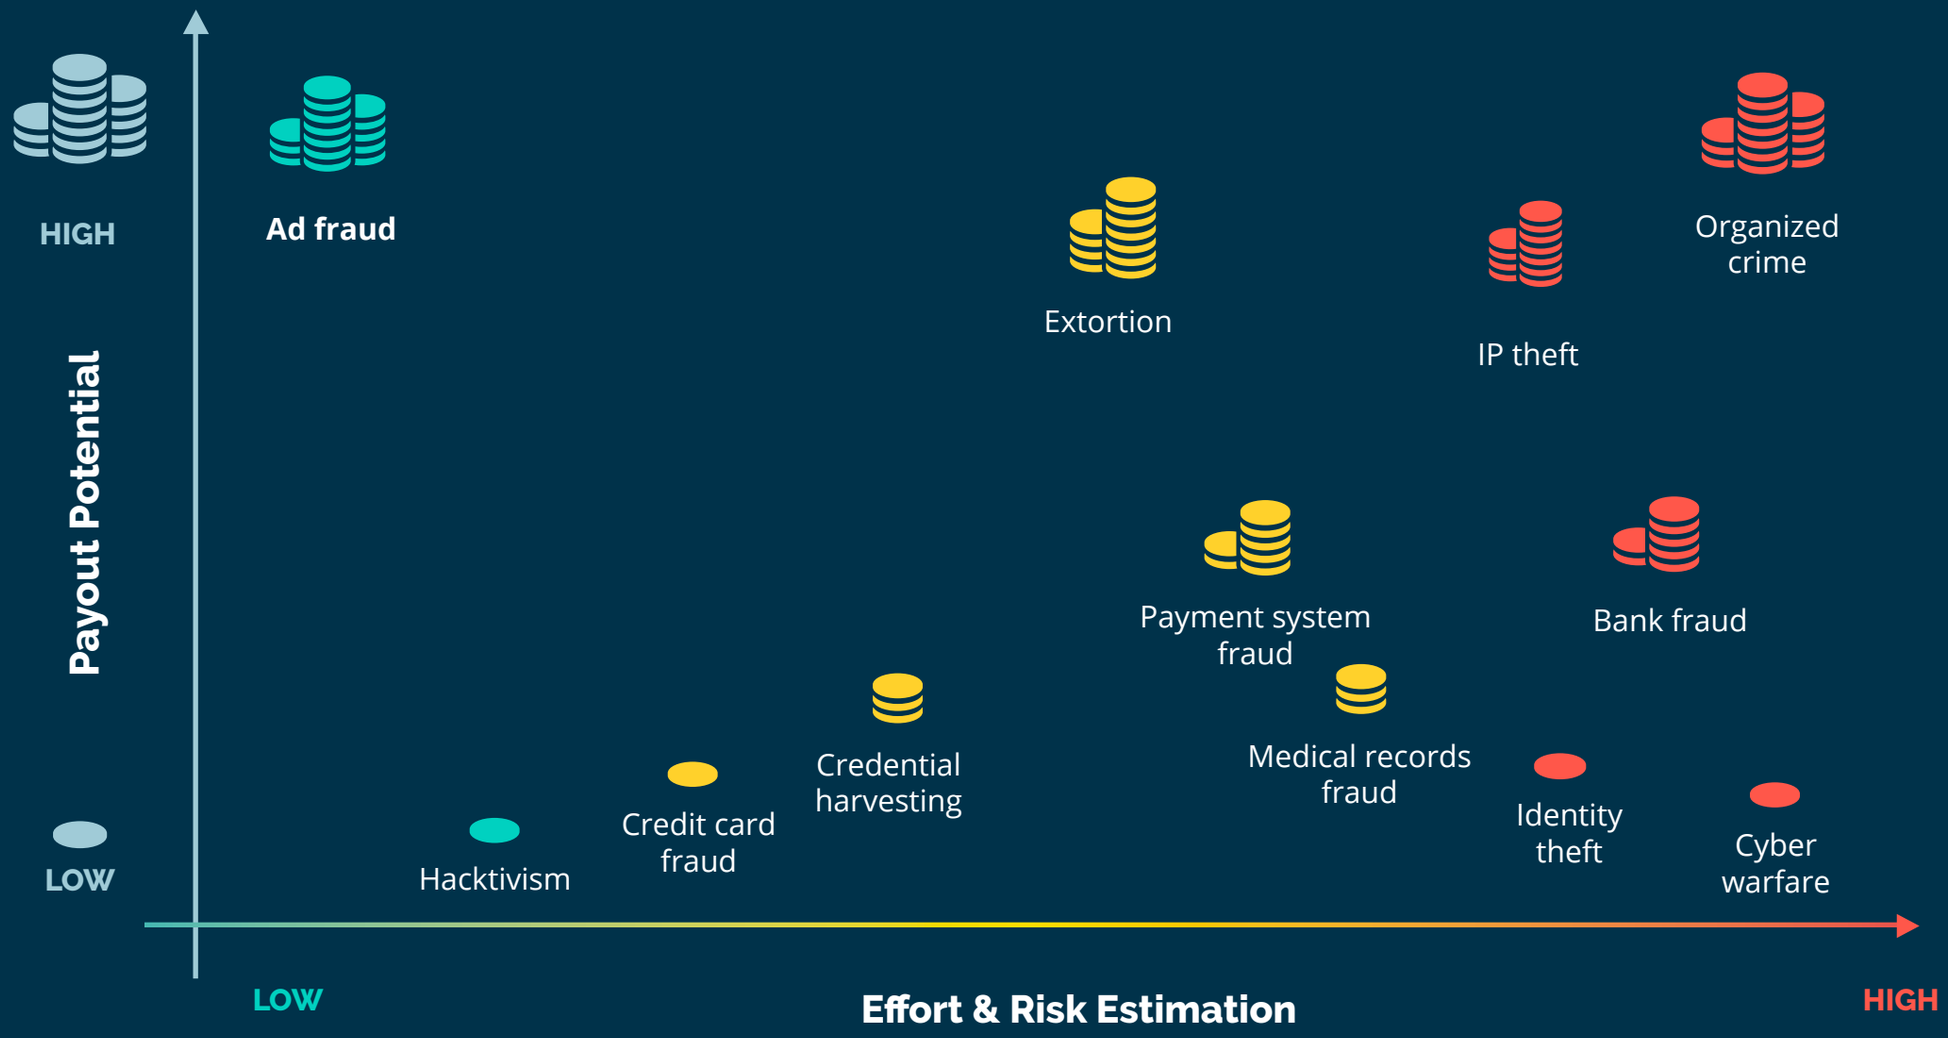

Source: Hewlett Packard Enterprises, “The Business of Hacking”, May 2016

Who's involved in perpetrating ad fraud?

The process of delivering a fraudulent impression involves multiple parties at different stages — starting from a publisher, then a network or exchange, then a traffic broker and/or malware distributor. The process is so interconnected that it is nearly impossible to determine who is at fault.

Ad fraud schemes are normally carried out by hackers and/or **bot network** (botnet) operators, but occasionally there are professionals who appear legitimate contributing as well.

Hackers and botnet operators

Hackers write the code to break into a computer and take control of it. Botnet operators coordinate with traffic brokers and issue instructions for the bots to follow.

Traffic broker

These middlemen connect websites looking to boost their traffic with the botnet operators who can supply it.

Other parties

Unfortunately, there are people with ties to the advertising industry who use technological manipulation and deceit to make money for themselves. They might create bogus content networks, funnel trade secrets to hackers, or even use a legitimate business as a front for selling fake traffic.

How do fraudsters get paid?

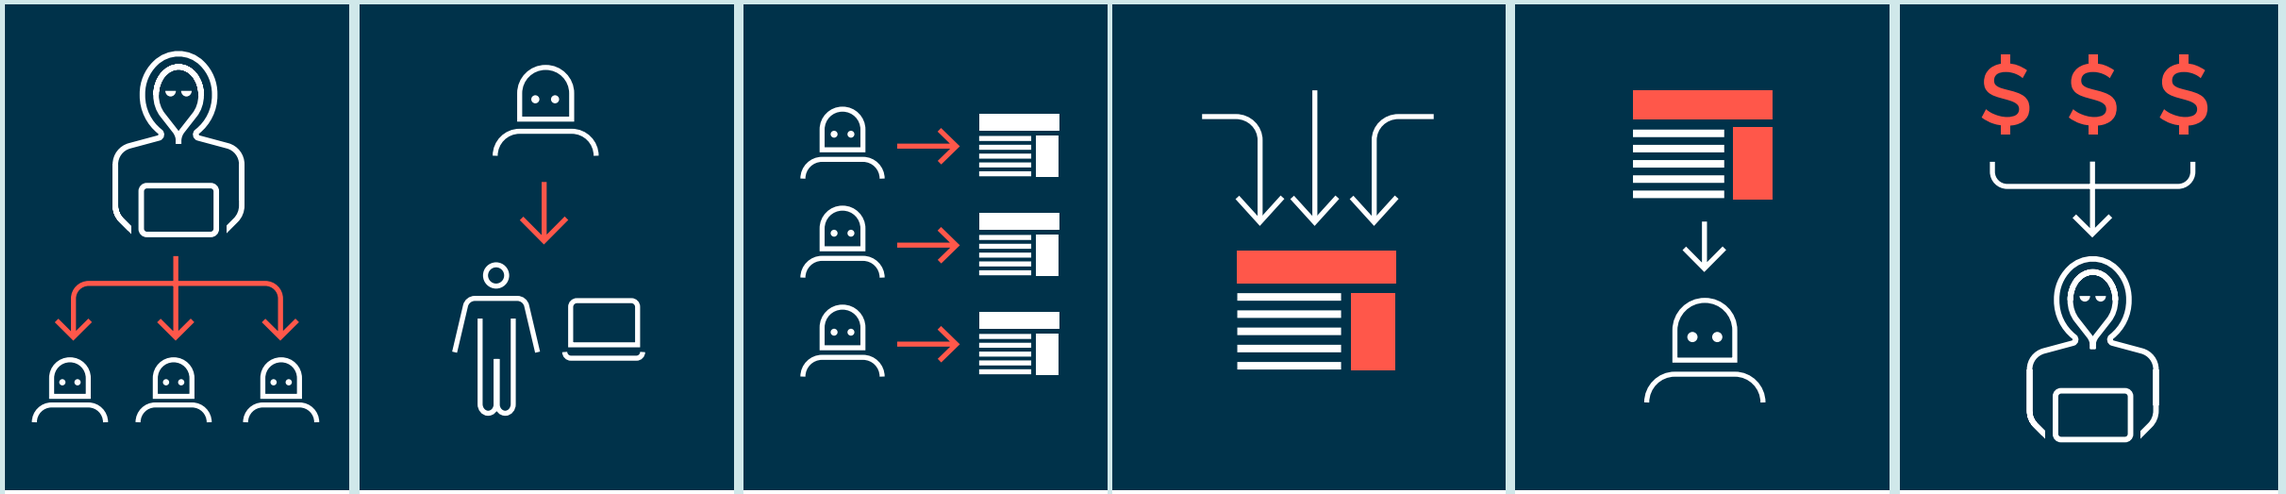

Hackers use code to create bots that then take orders from a botnet	Users unknowingly download and install bots on their computers	Bots are instructed to visit premium sites, picking up desirable cookies, and then visit fraudulent sites	Highly trafficked fraudulent sites use exchanges and networks to attract advertisers	Ads are continuously served to bots	Botnet operator gets paid

Key takeaways

Ad fraud is any deliberate activity that prevents the proper delivery of ads to the right people at the right time, in the right place.

The landscape of fraud is ever-changing, and every traffic source requires constant re-evaluation.

Ad fraud can be challenging to identify as it comes in many shapes and impacts each market differently.

The industry classifies fraud into two categories: General and Sophisticated Invalid Traffic.

CHAPTER TWO

Types of fraud

The aim of advertising is to deliver the right message to the right person in the right environment. Fraudsters use various techniques to compromise all of these three core values, resulting in wasted advertiser spend and damaged reputations for susceptible publishers.

IAS has a dedicated **Threat Lab** comprised of specialized analysts, engineers, white-hat-hackers, and data scientists who identify and stop fraudulent activity. Our team also works with experts from law enforcement and academic institutions to continuously examine the forces behind digital advertising fraud.

Bots

When most people think of ad fraud, they think of malicious bots.

Short for robots, bots are software programmed to carry out automated and repetitive online tasks, and can do them much faster than human users. Some bots are “good”, such as search engine bots that index content for search engines. Other bots are “bad”, and these are the most common example of ad fraud. Bots vary in levels of sophistication and structure and perpetrate ad fraud in slightly different ways.

Malicious bots are becoming more sophisticated; they can be programmed to intentionally view ads, watch videos, click on ads, and will be used as a tactic to siphon off money from advertising transactions.

Bot traffic is a useful tool for fraudsters as it is hard for the industry to identify who is behind this traffic.

Hackers create bots to surf the web, click on ads and play videos, which drives up traffic, resulting in more money paid out to the fraudsters.

These bots are viruses that can be installed unknowingly on a computer and then use computer resources in an unnoticeable way. Most people with infected computers are completely unaware.

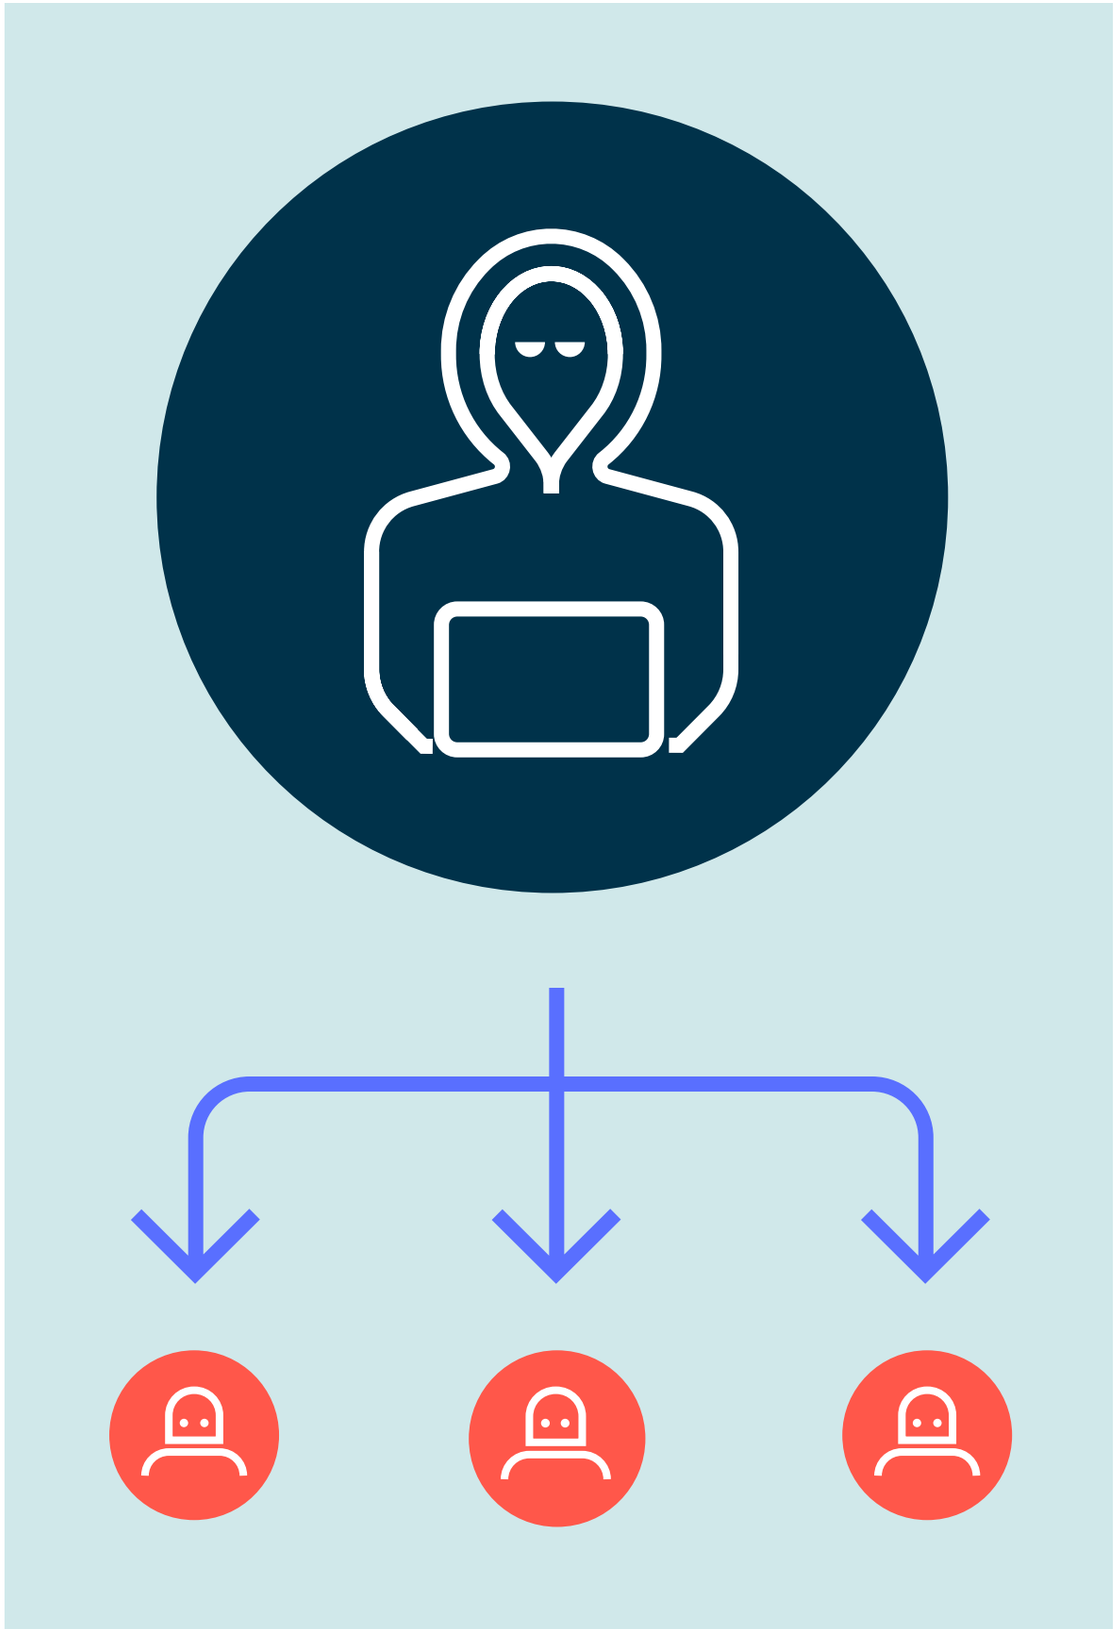

How does an infected computer become part of an illegal bot network?

A user's computer might be infected in a few different ways: by opening a phishing email attachment, clicking on malicious links, or installing software from untrustworthy sources.

Upon infection, the computer is given commands by the bot, making it part of a larger network of infected machines. The bot begins communicating with a server, which gives instructions for the bot to follow. These commands can include activities such as visiting premium websites in order to pick up cookies that typically define a desirable target audience to advertisers. The bot will then visit phony sites that buy traffic, creating more opportunities for ads to be wasted on bots.

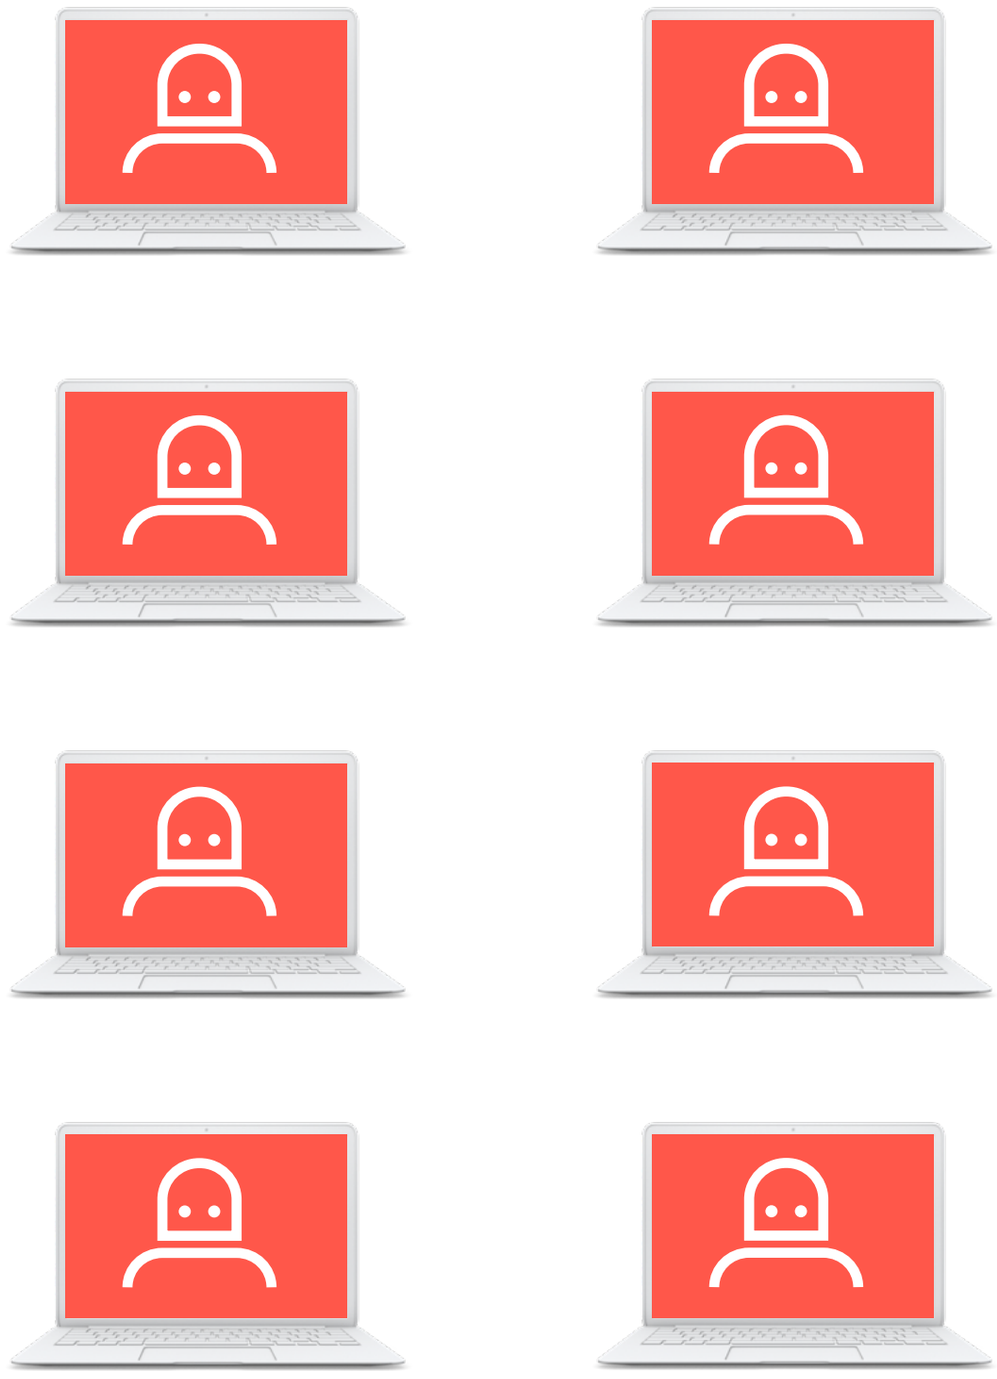

Types of bots

There is an increasing demand for transparency when it comes to ad fraud reporting, especially for **bot traffic**. Advertisers and media partners need to have more informative conversations about fraud in order to mitigate risk within campaigns, which requires more information in general.

There are many bots out there, acting more and more like humans. The IAS Threat Lab heavily invests in bot research to uncover the latest threats posed to advertising campaigns. Here's a quick look at some of the major bots IAS has uncovered, to give you a sense of how different bots can operate.

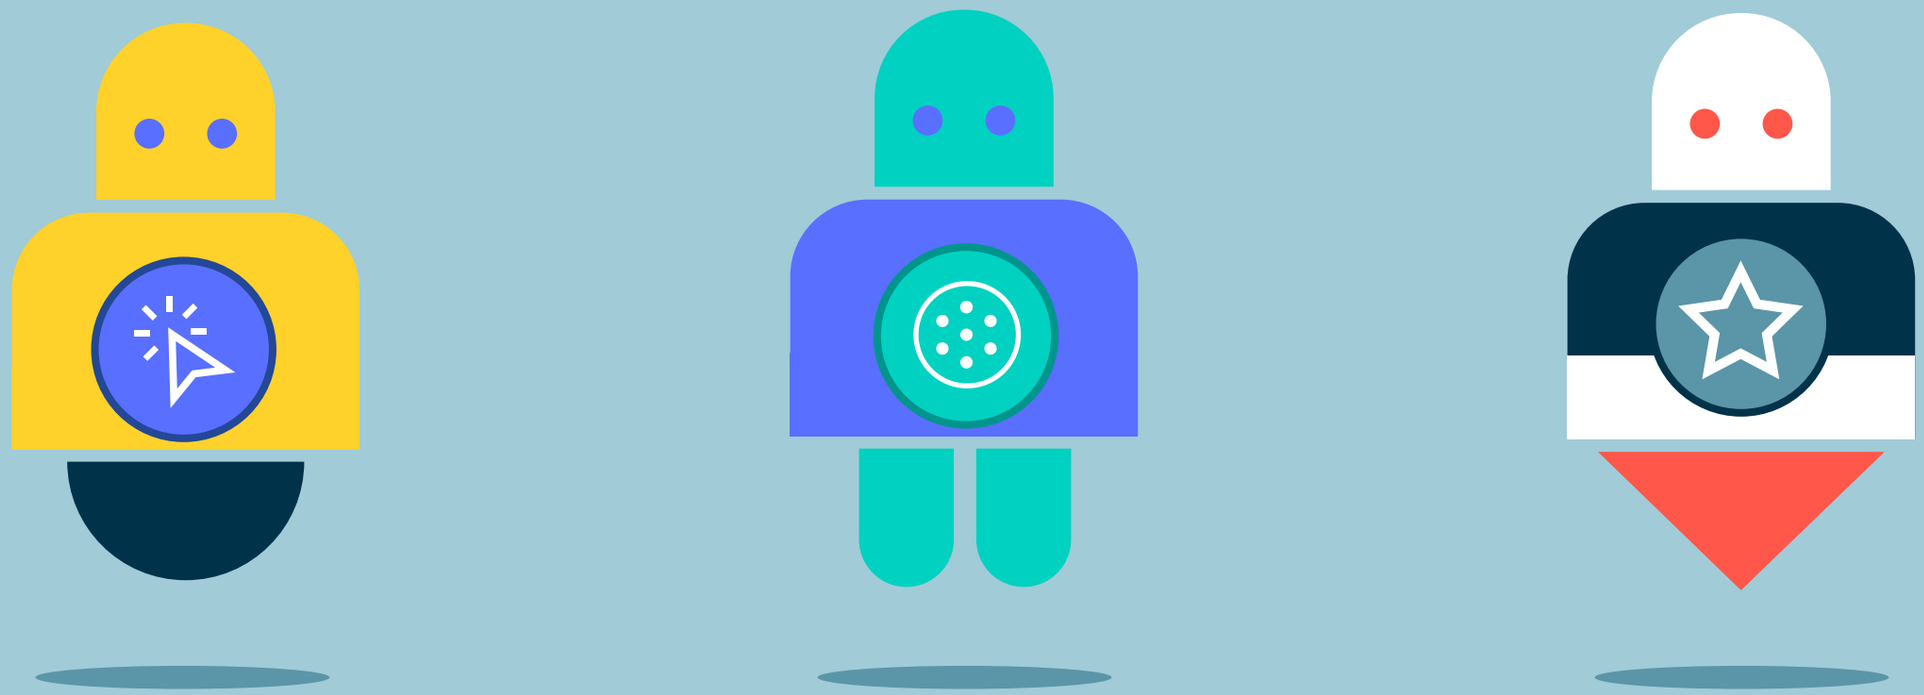

Poweliks

Poweliks is a botnet deployed to commit impression and click fraud by generating traffic to ad-supported websites and faking human interaction. It has the capability to execute up to 25 hidden browsers, all simultaneously conducting fraud. It has been specifically programmed to evade certain detection techniques used by ad tech companies.

Avireen

Avireen is a bot also used for conducting both impression and click fraud. It impersonates human behavior, simulating mouse movements and hovers. Avireen leverages the user's cookie cache to masquerade as a believable, real human user and deletes any new cookies that could flag it as a bot.

Proxy8

One of the most powerful bots IAS has detected to date. Within the first months of detection, Proxy8 compromised more than 60,000 machines around the world and earned an estimated $40,000 a day. With a working kernel-mode driver, Proxy8 has a level of power rarely seen in malware today.

Domain spoofing

Domain spoofing is a form of fraud where a fraudster impersonates a company's domain in order to pass off low quality inventory as high quality. Fraudsters fool buyers into thinking their ad is going to a premium site, when in reality it's going to a low-quality website. The impressions and the users are real, but the inventory is falsely represented and therefore purchased at a much higher cost.

Domain spoofing is also commonly used to mask unsafe sites. Brand safety is a huge concern to advertisers, and fraudsters take advantage by spoofing the domains of sites, like video piracy sites, in order to conceal their real identity and monetize the traffic.

Costs of domain spoofing:

» Loads programmatic buys with low-quality inventory

» Violates the security of whitelists

» Throws off KPIs

» Steals advertiser spend

» Steals publisher potential ad revenue

For more on domain spoofing, check out our IAS Insider post on the four types of domain spoofing.

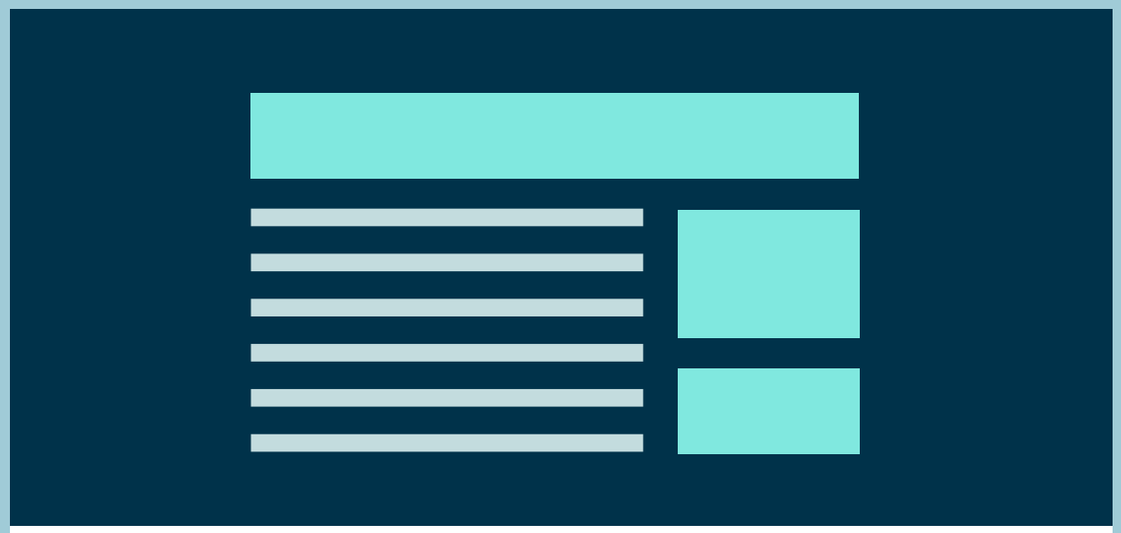

Sites that look real

Have content stolen from legitimate sites

Receive invalid traffic from any source

Might in fact be filled with ads which are hidden from a user's view

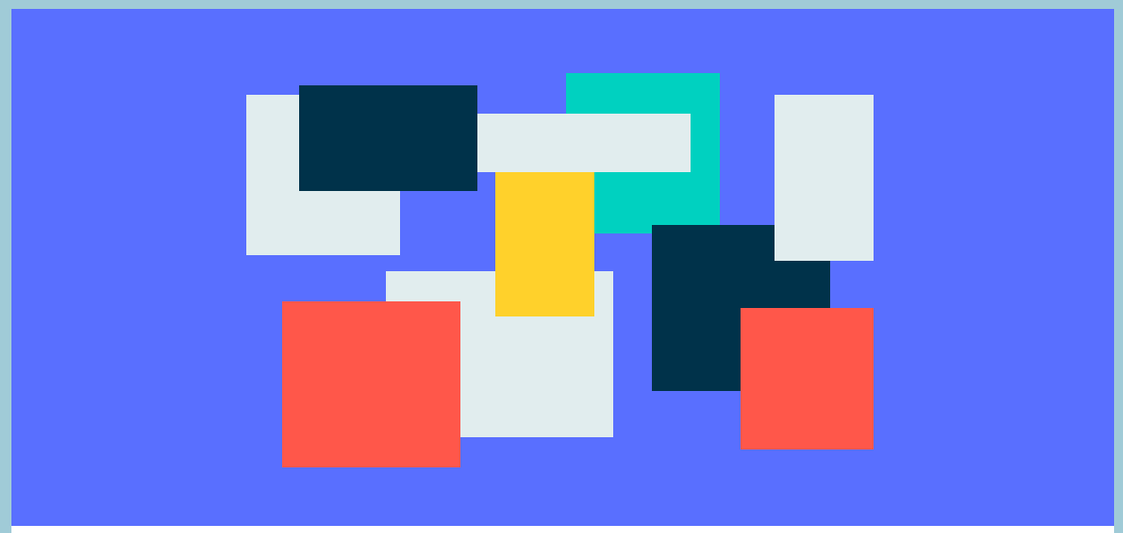

Sites that are only filled with ads

Cause a significant drain on a device's resources

Are designed primarily for visits by bots

Different methods of fraud can be combined for maximum effect, and any given impression can be fraudulent in more than one way.

For example, fraudsters often act as fake publishers, creating websites that contain ads. These sites might steal content from other pages and appear to be a normal page, or the sites can be solely ads without any intention to attract eyes.

Other forms of ad fraud

PIXEL STUFFING

Serving one or more ads, or an entire ad-supported site, in a single 1x1 pixel frame so ads are invisible to the naked eye.

AD STACKING

Placing multiple ads on top of each other in a single placement, with only the top ad being viewable. The advertiser is paying for impressions even if the user is not seeing the ads.

LOCATION FRAUD

Advertisers pay a premium for their ad to be served in a particular region, but fraudsters will send false location information so the ad actually serves elsewhere. For example, users might be surfing the web on their mobile device in New York City, and see ads for last minute tours of Alcatraz, California nearly 3,000 miles away.

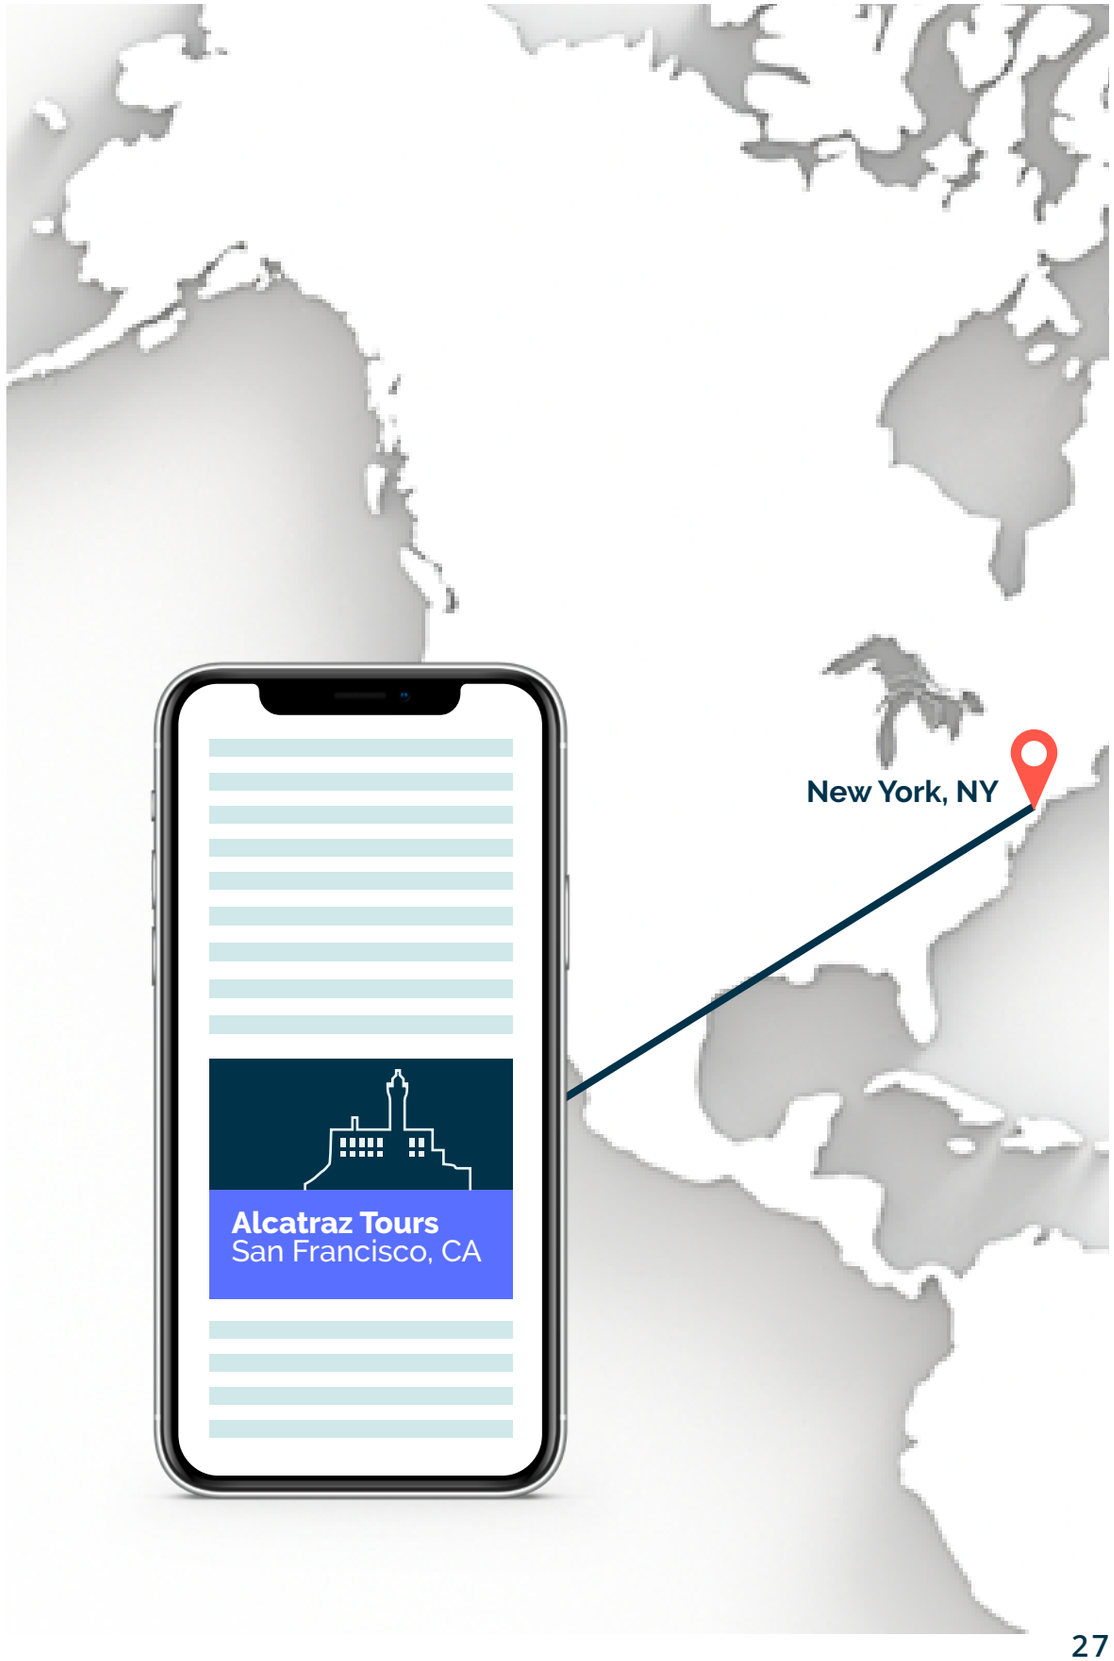

COOKIE STUFFING

Cookies are a method of tracking user behavior, to help determine what advertising effort led to a conversion (click, purchase, etc.) or what a user's interests are.

Cookie stuffing can happen in different ways. Fraudsters may try to game attribution models by adding a cookie to a user from an entirely different website from the one that the user originally visited. If the user later converts, the website associated with the stuffed cookie gets credit — and gets paid — for that action.

Cookie stuffing can also refer to the practice of placing many cookies on a bot so that they get targeted at higher CPMs, even if they haven't been flagged as potential high-value consumers.

USER-AGENT SPOOFING

Every request for a web page is sent with a “header” that provides some basic information about where the page is being loaded. One such piece of information is a description of the browser: its type, version, operating system, even plug-ins.

In user-agent spoofing, this description is modified to obfuscate information about the browser being used, which can interfere with user targeting. It’s most often used by bots trying to hide their tracks.

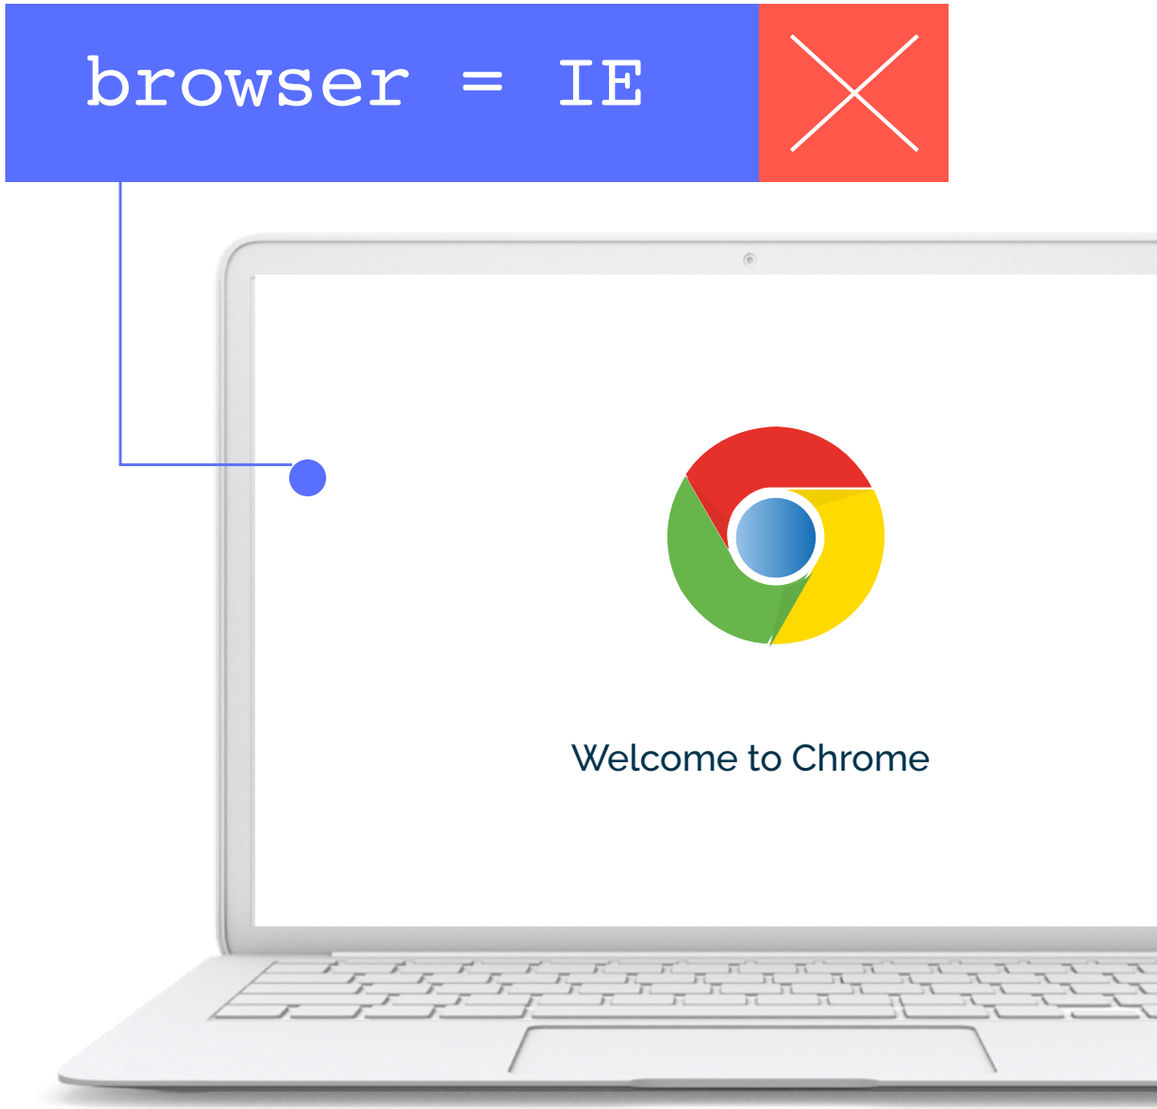

Key takeaways

The most prevalent forms of ad fraud are illegal bots and domain spoofing.

Bots are viruses that can be installed unknowingly on a computer and that use computer resources in an unnoticeable way.

Domain spoofing is when fraudsters fool buyers into thinking their ad is going to a premium site, when instead it's going to a low-quality website.

Different methods of ad fraud can be combined for maximum effect and maximum profit gain.

CHAPTER THREE

Fluctuating ad fraud

The impact of ad fraud can vary dramatically across different environments and formats. There can be fluctuations depending on your location and your purchase method or channel.

Ad fraud and video

Video consumption is becoming more fluid across screens, and digital video continues to attract greater consumer attention with audiences increasing worldwide. As more viewers flock to digital video, ad spending will also grow. Zenith forecasts that global display ad spending, which includes video and social media, will overtake linear TV as the largest component of ad spending by 2021.

Advertisers pay a premium to connect with consumers through sight, sound, and motion, and fraudsters know it. Common video ad fraud techniques involve misrepresenting display units as video placements to take advantage of higher CPM payouts.

Video inventory has been particularly susceptible to ad fraud because fraudsters understand that the medium with the greatest expected consumer impact can attract the highest CPMs. The video-heavy revenue models of many premium publishers can also make them more susceptible to fraud.

During the second half of 2018, we saw that globally, display video ad campaigns that had not taken any form of ad fraud mitigation saw nearly 1 in every 10 video ad impressions being impacted by fraud.

Global non-optimized fraud levels, H2 2018

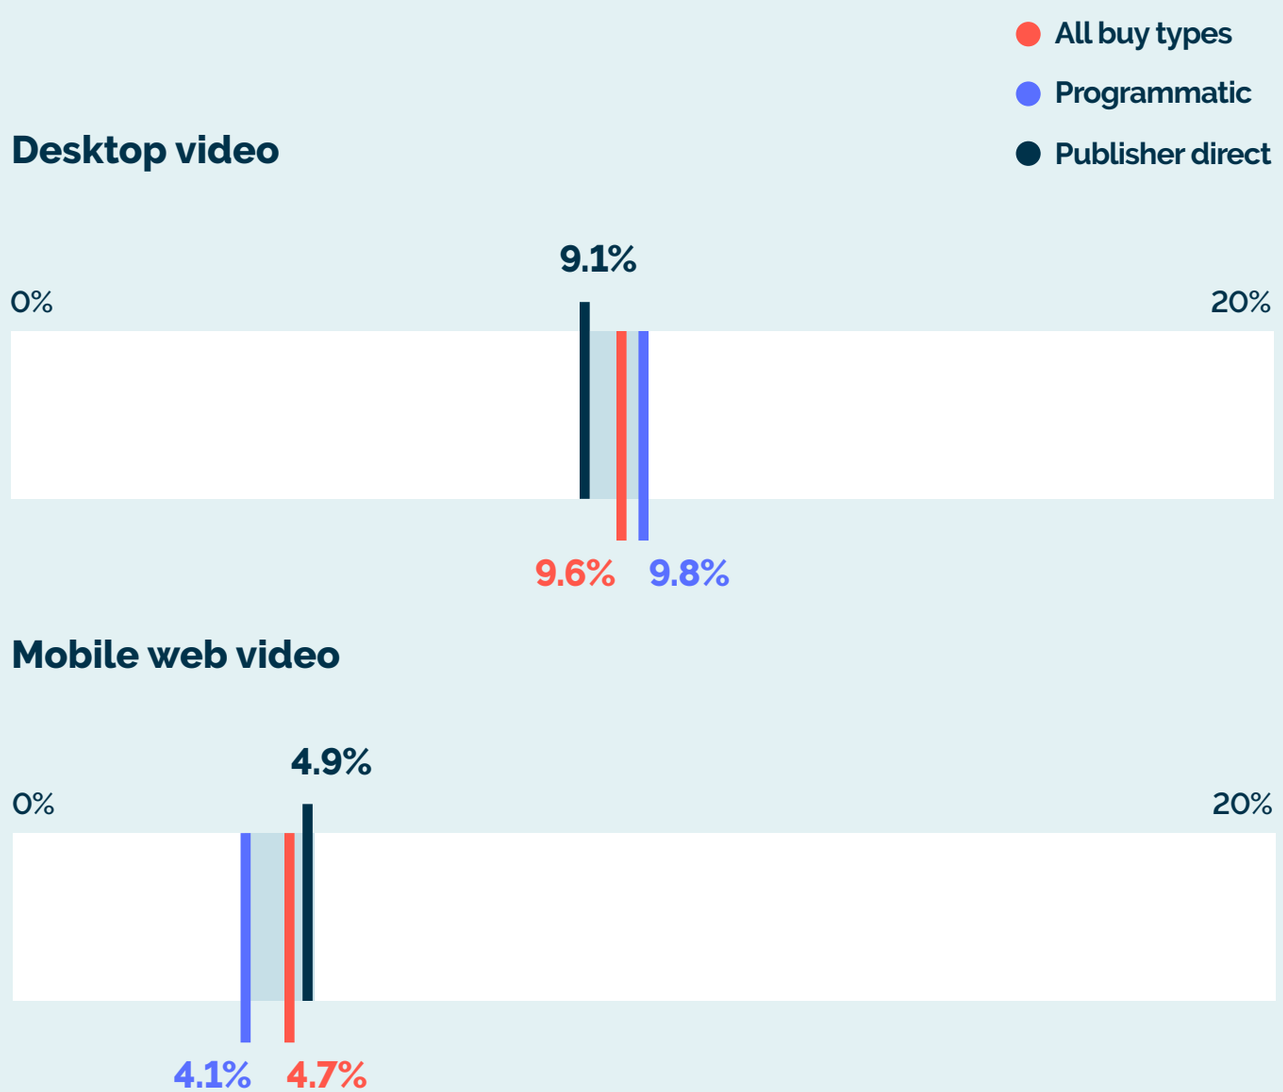

Other common forms of ad fraud in video

BANNER STUFFING

Banner stuffing happens when multiple ads are stuffed in the same ad slot, but only one is visible. It is a form of pixel stuffing.

VOLUNTEER BOTNETS

Everyday users knowingly allow their computers to be infected by bots that will automatically browse content, in order to get reciprocal traffic to their own content. Volunteer botnets often focus on individual videos within user-generated content sites so that content uploaders can share in the revenue.

Ad fraud and mobile

Approximately 52% of the world's population — 3.9 billion people — access the internet via mobile. In 2018 alone, advertisers spent an estimated $137.9 billion to reach mobile users. Mobile is currently poised to overtake television as the largest ad format.

As mobile continues to grow in consumer usage — and as advertising follows — fraud techniques will become more pervasive and tailored to mobile.

Mobile ad fraud is becoming increasingly sophisticated and more challenging to detect. When compared to other channels, mobile environments are found to be very differently impacted by ad fraud. Particularly the in-app environment, where display fraud deployments cannot be transferred over.

Global non-optimized fraud levels, H2 2018

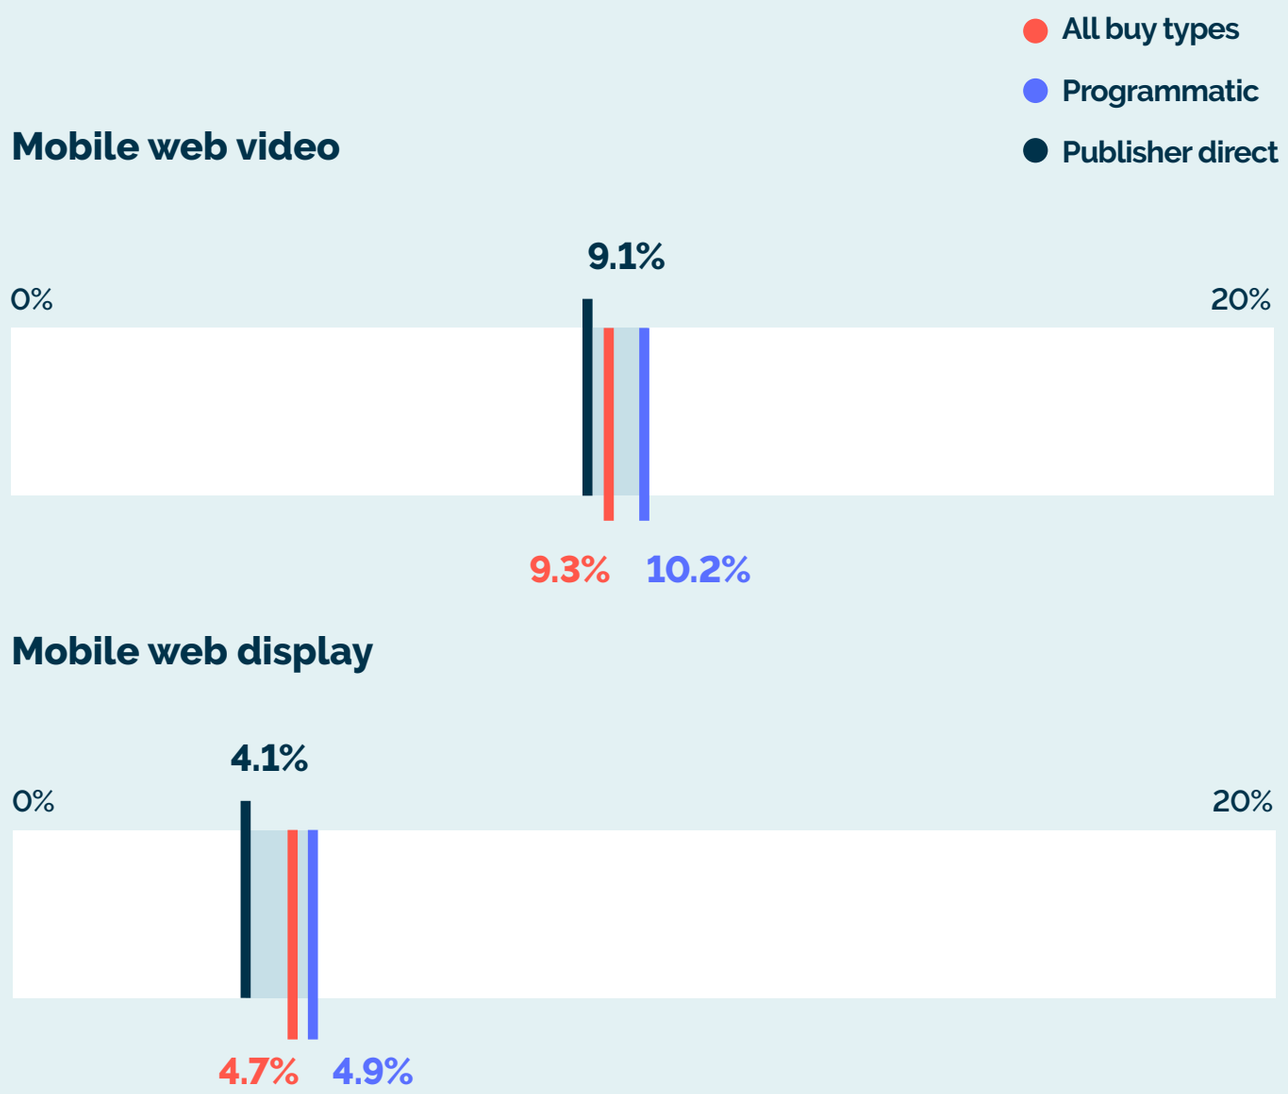

How mobile ad fraud occurs

While bots can still wreak havoc on the mobile platform, the more prominent type of mobile fraud involves hiding ads in services or apps running constantly in the background. Location spoofing and app-name spoofing are other costly forms of mobile fraud.

MALICIOUS APPS

Apps that generate fraudulent impressions without the user knowing. This can be thought of as a kind of mobile malware.

HIDDEN ADS

Common also in desktop fraud, hidden ads are generated in-app in a way that is not visible to the user.

Examples include:

» Ad stacking

» Banner stuffing

» Invisible banners

BACKGROUND SERVICES

Apps running in the background are able to render ads. These ads are invisible to the user and are capable of generating large ad volumes. Background service fraud commonly has the following characteristics:

» The app does not need to be manually started by the user. It auto-starts on every reboot or is triggered by a common activity, such as screen lock.

» Spoofs geo-location/device ID/app name in order to obscure abnormally high ad volume and remain inconspicuous

APP INSTALL FRAUD

If advertisers pay per install, fraudsters will find a way to falsify the number of installations.

This scam can be carried out either by fake mobile devices installing apps and getting credit for the install, or by thousands of real phones being commanded to download apps and generate paid install fees. Fraudsters can then reset a phone's ID to appear as a new device, enabling them to commit the fraud over and over.

SDK SPOOFING

Fraudsters have learned how to hack **software development kits (SDKs)** and manipulate advertising by mimicking user behavior. Any genuine app developer that integrates with the wrong kind of SDK is vulnerable. This type of fraud is particularly hard to detect because it occurs on real phones using legitimate apps. Malware hides on an app generating a series of simulated ad clicks, installs, and engagement signals to an attribution provider without any real installs occurring.

CLOUD HOSTING

In-app impressions are displayed on devices hosted in the cloud, generating ad revenue for the app creators. Fraudsters are able to control and change signals such as device ID and geo-location, making it appear as though there are a variety of different devices and therefore users, while in reality there is one hijacked device in the cloud.

APP NAME SPOOFING

Similar to domain spoofing in display, apps can submit a false app identifier to the bidding platform. This interferes with detection of apps utilizing background services to load ads, as well as brand safety and contextual targeting.

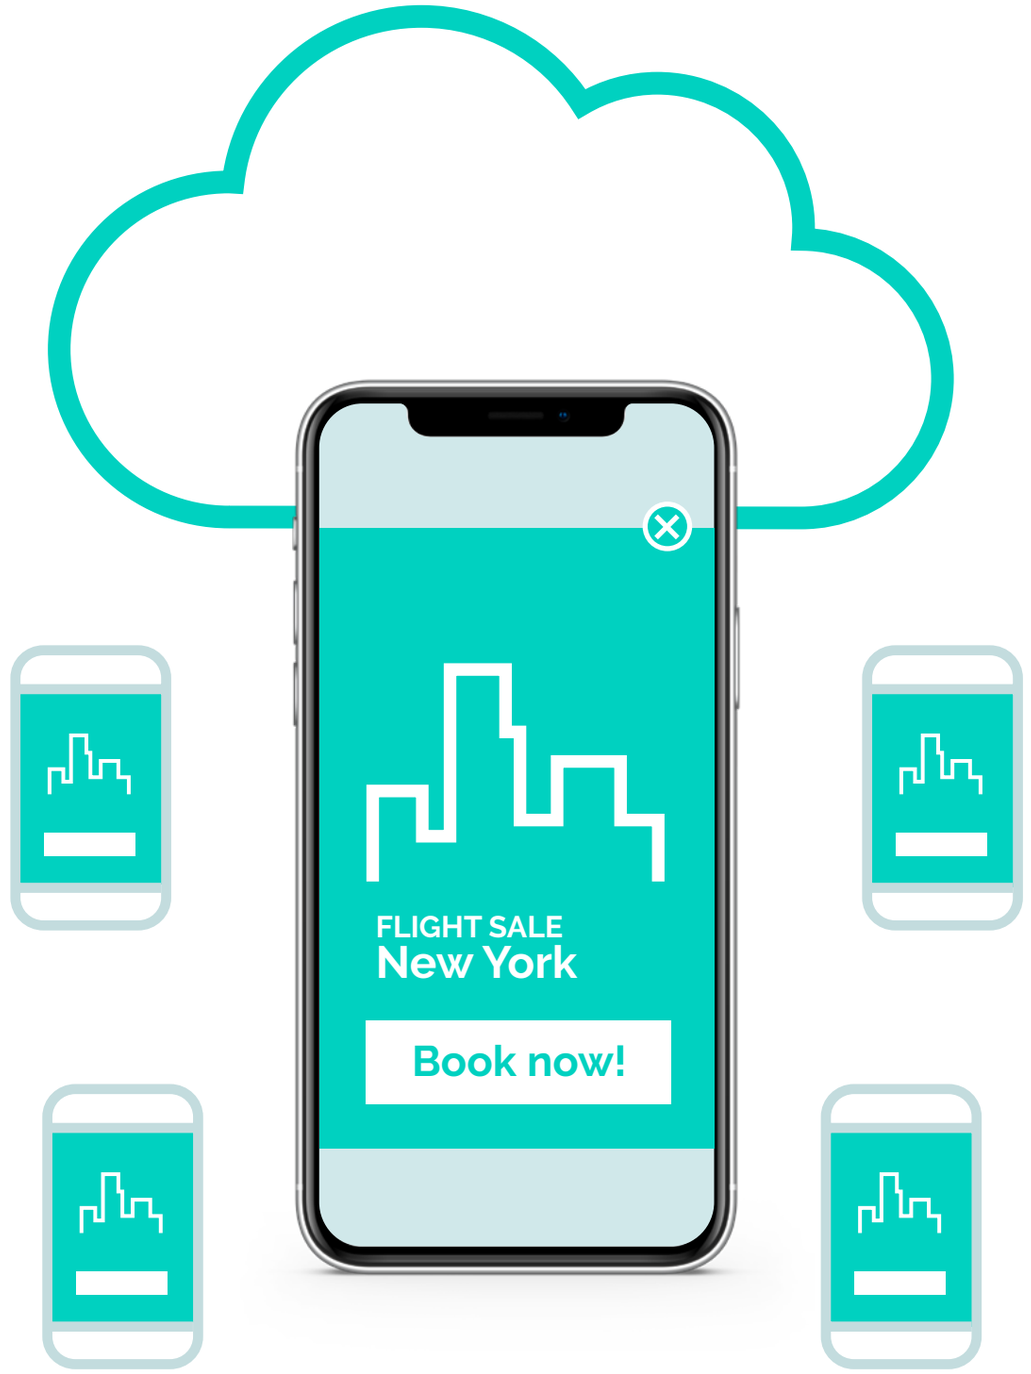

Why is mobile ad fraud particularly challenging?

Much of fraud in mobile advertising behaves similarly to that on desktop, but the technologies involved are entirely different. So, detection techniques have to be engineered from scratch, making it a particular challenge to combat fraud in mobile.

As technology continues to develop, and the **Open Measurement SDK** continues to be implemented, aligning and improving overall mobile measurement will help to reduce the impact of ad fraud in mobile environments.

Find out more about the Open Measurement SDK here and how it is enabling the ad industry's measurement of mobile.

Ad fraud in newer technologies

As advertisers direct more budget towards newer technologies, they need to keep ad fraud concerns front and center. With new inventory sources, demand is strong but the supply of impressions can be limited. When advertisers create more demand than the legitimate supply can meet, fraudsters take the opportunity to trick advertisers into buying inventory that does not really exist.

AD FRAUD IN CONNECTED TV

It's still early days for **connected TV (CTV)** and it is expected that this emerging technology will come with new types of ad fraud. It is clear however, that certain types of fraud are not possible in this environment. For example, cookies and clicks are not supported in a CTV environment.

CTV fraud scenarios include bots, device spoofing (where devices are obfuscating their attributes to maraude as legitimate), browser spoofing, and malicious CTV apps.

IAS worked with Verizon and eight leading video publishers to develop an industry-first solution to verify video ads that are played on fraud-free connected TVs. IAS verified hundreds of millions of ads and saw lower levels of fraud when compared to other digital environments.

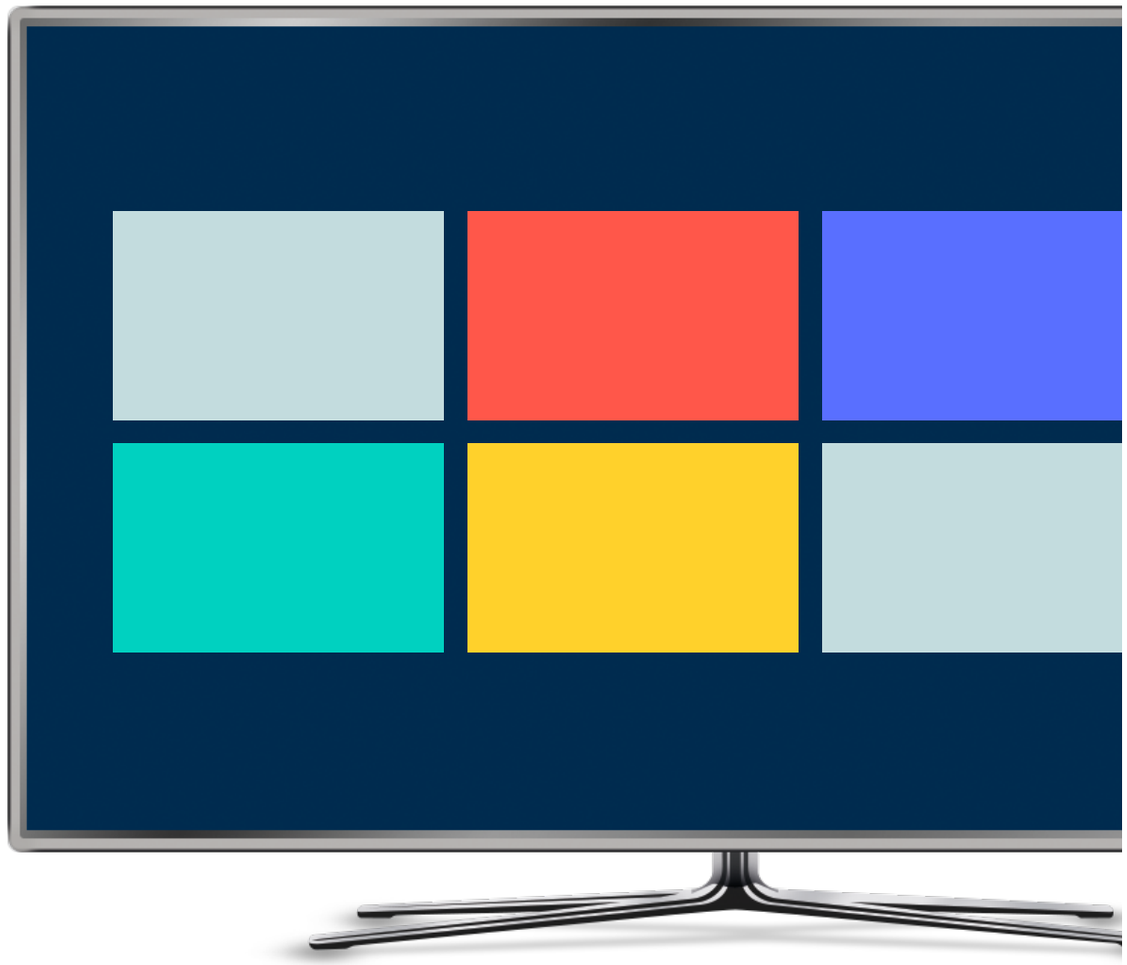

AD FRAUD IN SOCIAL

Much like in other environments, malicious bots are used to commit the majority of ad fraud on social media platforms. These platforms are not immune to compromised computers with breached security defenses conceded to a third party. In this scenario, bots will carry out automated actions in the background of the social media platforms without the user knowing.

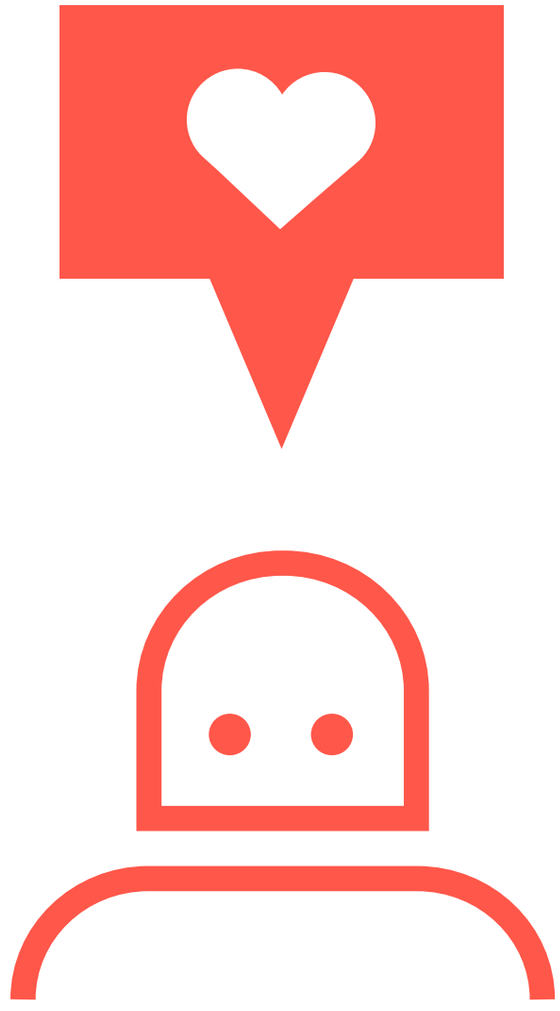

Key takeaways

Ad fraud fluctuates and can impact every part of a media plan.

As digital video continues to grow, not only will it attract greater consumer attention but it will also attract frausters' attention.

As mobile use has risen so has mobile advertising and as such, ad fraud techniques will become more tailored to mobile.

Newer technologies are just as susceptible to fraud and will be the next focus for fraudsters.

CHAPTER FOUR

How to fight ad fraud

As ad fraud evolves, the digital industry needs more sophisticated methods of detection. We know that the fraudscape evolves quickly, and that fraudsters are always finding new ways to evade detection, making the situation ever more complicated.

Furthermore, not all aspects of ad fraud are explicitly illegal, and those that are usually occur in countries with ineffective cybercrime law enforcement. With criminal justice systems lagging behind, proactive measures for avoiding fraud are required,

IAS takes a multifaceted, three pillar approach to ensure the most precise fraud detection possible.

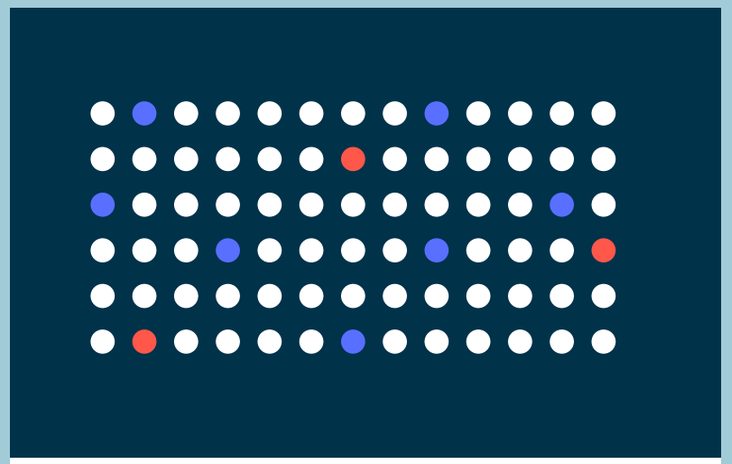

Behavioral and network analysis

(Recognizing differences in behavior patterns)

Massive data sets are collected and used to distinguish real human behavior from bots and identify anomalies indicative of fraud. This “big data” is then used to create accurate and scalable detection models.

Browser and device analysis

(Validating that the browser is real)

Each bot has a signature set of characteristics that can be identified through detailed mapping of the browser environment and device characteristics. A machine infected with a bot will perform differently than an uncompromised computer, despite the fact that the abnormality may not be apparent to the user.

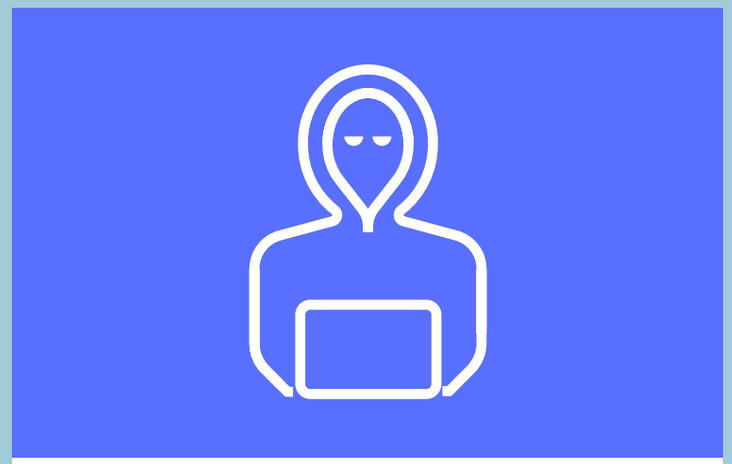

Targeted reconnaissance

(Dissection of malware and infiltration of hacker communities)

The IAS Threat Lab is increasingly engaging with the cybersecurity community to determine the most effective ways to infiltrate hacker communities and discover threats to advertisers. Just as fraudsters try to reverse engineer security signals from tech companies, malware analysts can reverse engineer bots.

How can the industry fight invalid traffic?

As combatting ad fraud has become a major industry priority, a range of industry bodies and initiatives have been established specifically to tackle the issue.

The **Interactive Advertising Bureau (IAB)** provides guidelines for the detection and filtration of invalid traffic to establish minimum requirements for identifying and removing ad fraud from advertising transactions.

Additionally, the **Trustworthy Accountability Group (TAG)**, a joint industry initiative, created the Certified Against Fraud Program to combat invalid traffic.

THE SUPPLY CHAIN CLEANUP

The nature of programmatic makes it easier to conceal the supply chain as well as the quality of the inventory, leaving ample opportunity for fraud. In an effort to address this "murky supply chain", **Ads.txt** was created. Publishers display a simple text file on the Ads.txt site that lists every company authorized to sell their inventory, giving both publishers and advertisers the ability to cross-check sellers and inventory and eliminate any opportunity to purchase fraud.

Ads.txt was a positive step towards identifying ways for publishers and advertisers to address fraud. However, IAS views Ads.txt as a tactic that protects against human-driven behavior, not bot-driven behavior.

For publishers, it is important to diligently manage what they put in their Ads.txt files, but also to audit their monetization partners closely. As a result, exchanges will feel more of an incentive to accurately screen against publisher Ads.txt lists and ensure they are not partaking in ad fraud.

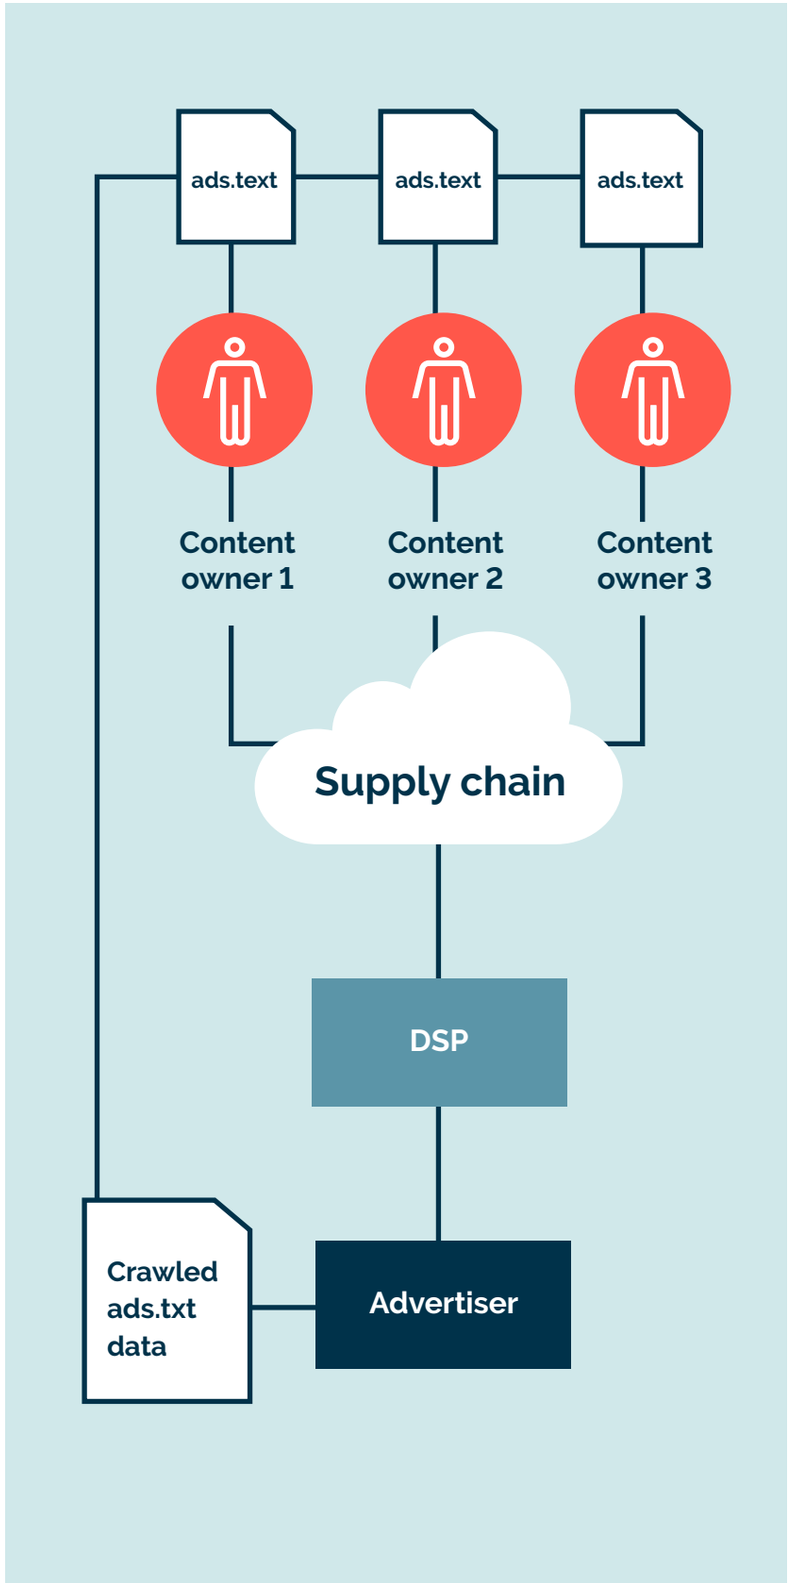

EXCLUSION LISTS

Advertisers who rely on exclusion lists apply a reactive method that immediately shuts down a supply channel as soon as fraud is detected. However, it is often the case that lists are not updated frequently enough, and using them can significantly impact scale. Even premium publishers who dilligently follow best practices can fall victim to fraud, and placing them on exclusion lists unfairly penalizes them.

Meanwhile, as soon as one site is placed on an exclusion list, a new one can be registered and used to continue the fraud, realtively uninterrupted.

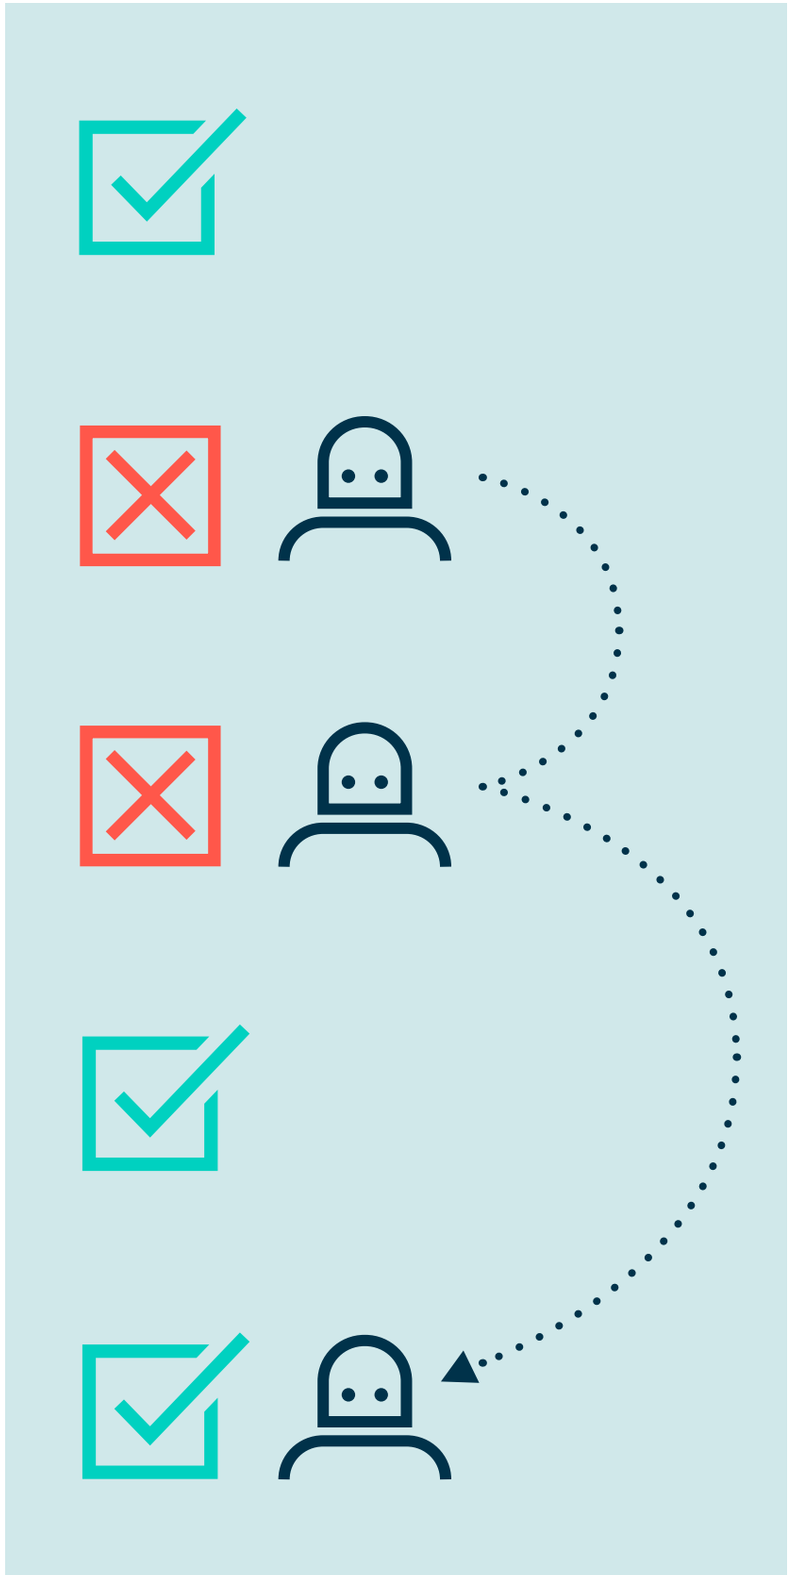

How to best protect yourself from ad fraud

The only real way to protect your campaigns is to be proactive in identifying fraudulent behavior and preventing it from impacting your campaigns in the first place. No single method is sufficient — you need combined, unified defenses to thwart ad fraud.

For publishers, vigilance is key. Relentlessly monitor inventory, **sourced traffic**, and vendors for ad fraud. Cut out anyone that's supplying bots. Use real-time measurement tools that can help you identify fraud the moment it hits your inventory. By enabling measurement across all inventory, you are helping advertisers monitor their media investments more closely, so you'll earn their trust and budget.

METHODOLOGY

The most robust solutions to fight ad fraud are those that combine methodologies. A layered approach to detection is vital, combining rules-based methods that evaluate at the impression level, along with machine learning that can identify likely fraud patterns.

Rules-based only

Reliant on manual adjustments to brand safety models which will not keep pace past today's cirital junction

Machine cognition

Multi-layer approach combines predefined rules with AI and Machine Learningin order to scale with the rate of content creation

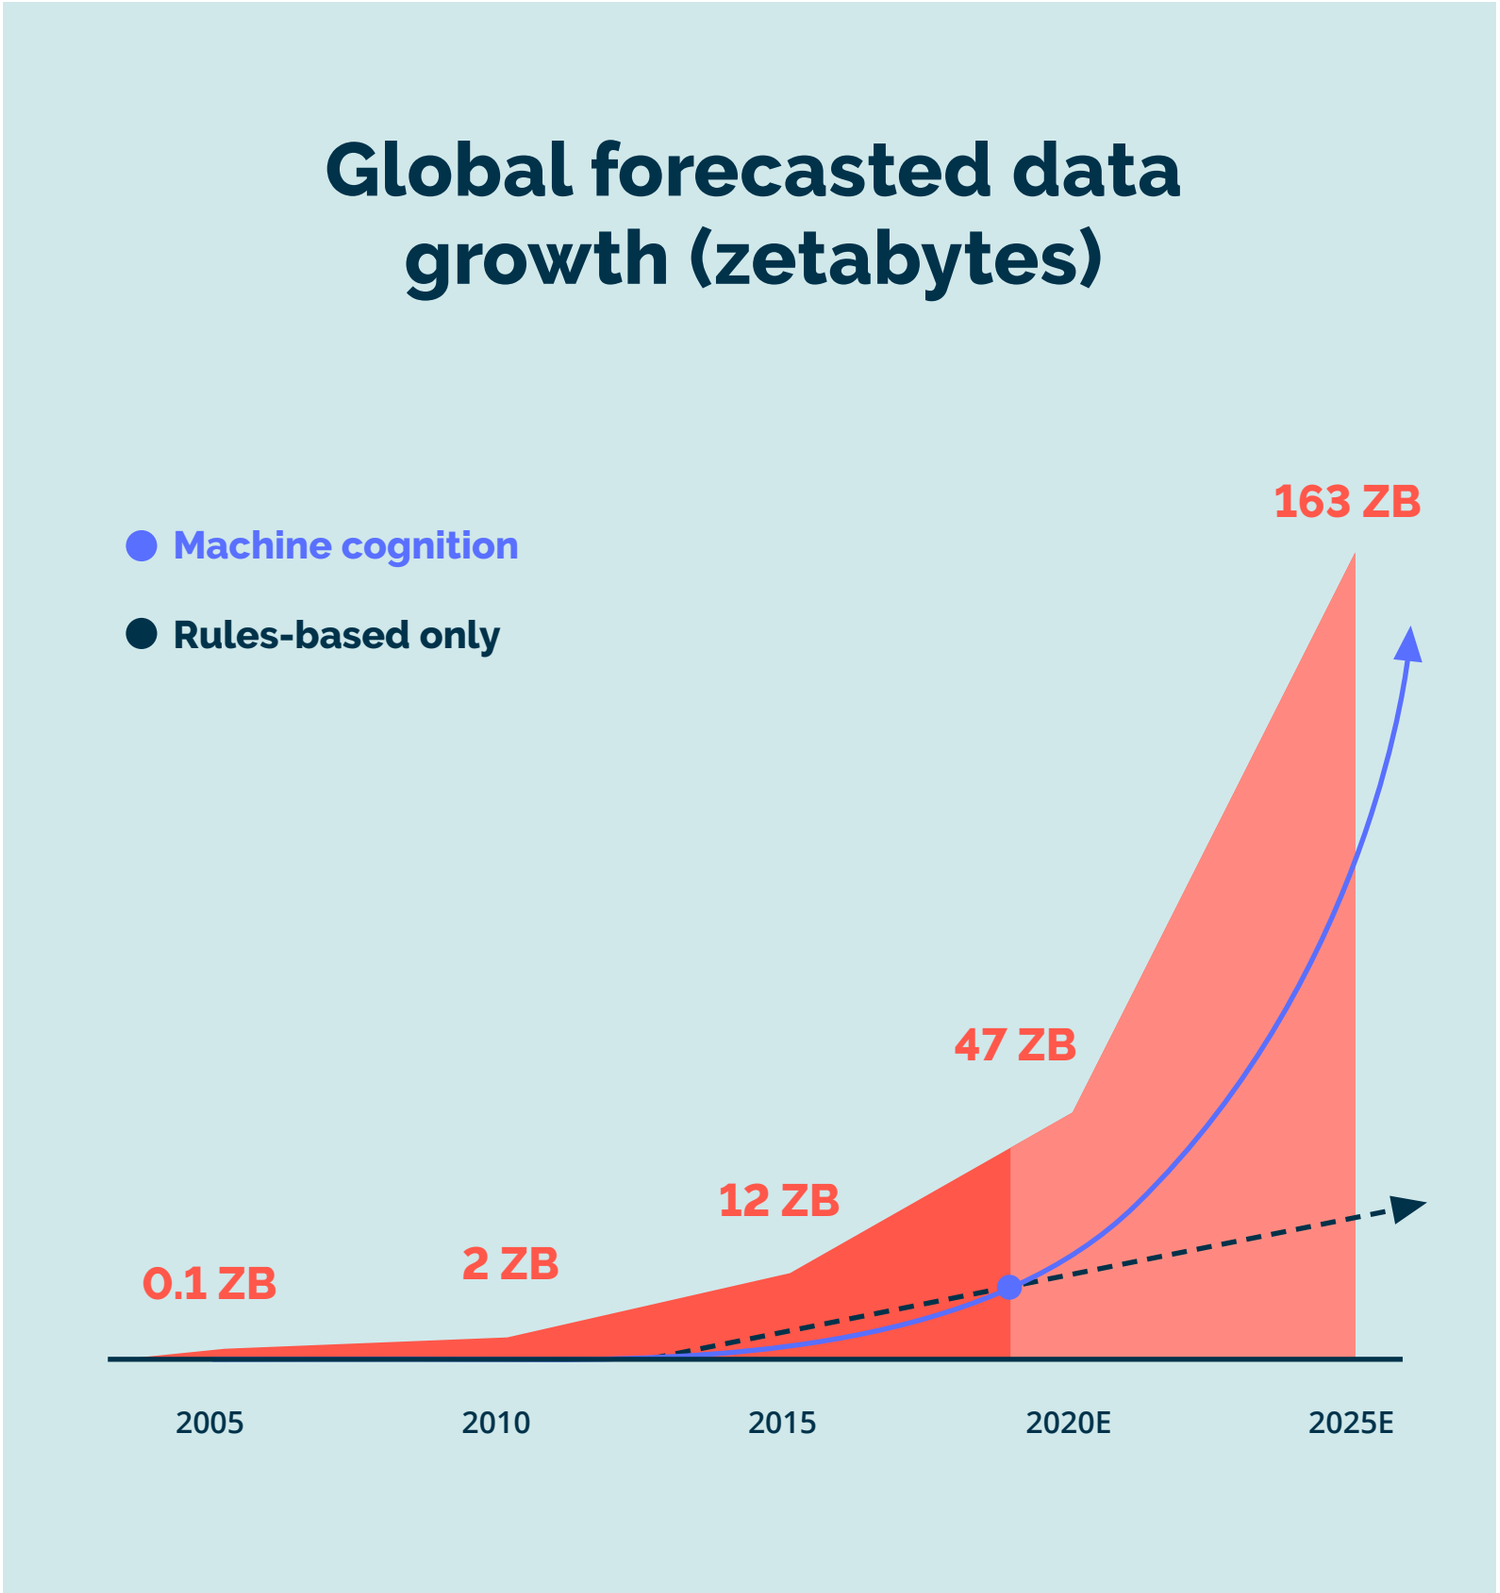

The biggest factor to protect yourself is awareness and participation. It is critical that all members of the digital ecosystem are a part of the effort to tackle ad fraud.

TEN TOP TIPS
TO PROACTIVELY PREVENT FRAUD

1. Use verification and fraud solutions that can confirm ads were delivered to the desired sites, devices, and geographies.
2. Use fraud solutions that have been accredited for both General and Sophisticated IVT.
3. Measure fraud across all campaigns to understand anti-fraud performance.
4. Vet your vendors and partners; ask how they measure for malicious bots and other forms of IVT.
5. Offer and request more transparency into inventory and traffic, including sourced traffic and audience extension.
6. Implement blocking technology or use anti-targeting technology to avoid infected machines or pages with historically high fraud levels.
7. Use pre-bid filtering to avoid fraud in programmatic buys.
8. Where needed, use exclusion lists.
9. If it's too good to be true, it probably is. Focus less on low CPMs and more on hitting real KPIs tailored to your campaign goals.
10. Offer and request more visibility and transparency into where programmatic advertising is being served.

CONSIDERATIONS FOR THE SELL SIDE

Ad traffic may be deemed invalid as a result of non-human traffic or activity designed to produce fraudulent traffic. Bots can be given instructions to visit premium sites, pick up desirable cookies, and then visit fraudulent sites to monetize. If you're selling premium inventory, this may explain some level of fraud on your website.

Advertisers' money is on the line, but so is the sell-side reputation and potential revenue. Here are the top four questions to ask yourself to best fend off undesired ad traffic and keep advertisers spending with you:

1. Where is your traffic originating from?
2. In what environment is fraud occurring the most (e.g. mobile web, desktop)?
3. Do you use an anti-malvertising solution that re-plays impressions and linked tags?
4. Do you source traffic or have a lot of data center traffic? Are you accredited for general and sophisticated IVT detection?

Read what to do if you discover invalid traffic and consult our checklist to fend off undesired traffic

Key takeaways

As fraud techniques advance, the digital industry needs more sophisticated fraud detection.

A layered approach to detection is vital and the most robust solutions to fight ad fraud are those that can combine technologically sophisticated methodologies.

Proactive measures to avoid fraud are required, and methods such as exclusion lists can significantly impact scale.

IAS takes a three pillar approach to ensure the most precise detection and prevention of ad fraud.

CONCLUSION

Ad fraud is everybody's problem.

Digital advertising is constantly evolving to achieve marketers' goal of delivering the right ad, to the right person, at the right time. This is made apparent with a projected $500 billion to be spent on digital in the near future. At the same time, we have discovered the world of invalid traffic is growing just as quickly. As new innovations like CTV and other technologies appear on the horizon, so too will new advancements in ad fraud.

Just as every traffic source requires constant re-evaluation, so does industry education on the detection and removal of ad fraud. In order to combat fraud, we all have to take on the responsibility. This means understanding what the issue is and how you can be impacted, alongside getting certified for GIVT and SIVT filtration, and demanding your partners do the same. It means supporting cooperative industry initiatives in order to protect ourselves and our consumers as digital evolves.

Overall ad fraud takeaways

Ad fraud is defined as any deliberate activity that prevents the proper delivery of ads to the right people, at the right time, in the right place.

Different methods of ad fraud can be combined for maximum effect and maximum profit gain; it is imperative to understand how fraud fluctuates.

Ad fraud can impact every part of a media plan. Newer technologies are just as susceptible as older ones and will be the next focus for fraudsters.

As ad fraud becomes more advanced, the digital industry needs sophisticated detection methods and proactive measures to effectively avoid it.

Glossary

Ad fraud: Any deliberate activity that prevents the proper delivery of ads to the right people at the right time, in the right place.

Ad stacking: Placing multiple ads on top of each other in a single placement, with only the top ad being viewable. The advertiser is paying for impressions even if the end user is not seeing an ad.

Ads.txt: Created in an effort to help the industry identify fraud by allowing publishers to display a simple text file on its site to list every company authorized to sell their inventory.

App install fraud: If advertisers pay per install, fraudsters will find a way to falsify the number of installations.

App name spoofing: apps can submit a false app identifier to the bidding platform, interfering with detection of apps utilizing background services to load ads as well as brand safety and contextual targeting.

Banner stuffing: Form of pixel stuffing; multiple ads are stuffed in the same ad slot, but only one is visible.

Bot: a software program that carries out automated tasks on the Internet. There are good bots and bad bots.

Botnet: (Bot network) A group of computers taken over by malicious bots.

Bot traffic: Invalid traffic designed to mimic users and inflate audience numbers.

Cloud hosting: In-app impressions that are generated by hosting devices in the cloud and running apps that display ads, generating ad revenue for the app creators

Connected TV (CTV): A television that is connected to the internet, unlike traditional TV sets

Cookie stuffing: Process in which fraudsters add cookies from a different site than the consumer is actually visiting, in order to profit from the false action

Datacenter traffic: Traffic originating from servers in data centers, rather than residential or corporate networks. Typically, no end user is present, though proxy servers or other technologies may result in traffic appearing to originate from data centers while still being delivered to human users.

Domain spoofing: A form of fraud where a fraudster appears to use a company's domain to impersonate the company in order to pass off low quality inventory as high quality.

Exclusion list: Using lists of known bad IPs, domains, or other parameters to prevent the serving of ads matching those parameters.

General Invalid Traffic (GIVT): Traffic that comes from known, non-human sources on publicly available IP lists. This fraud can be identified through routine means of filtration.

Hidden ad impressions: Impressions that are not seen by users because they are hidden behind other ads or content, displayed in tiny iFrames (pixel stuffing), or otherwise served in a way that prevents real ad views.

Hijacked device: A user's device (browser, phone, app) is modified to request ads that are not under the control of a user and made without the user's consent.

IAS Threat Lab: A team of dedicated experts at IAS using tools to identify and stop fraudulent activity. Read more about the threat lab here.

Incentivized browsing: A human user may be offered payment or benefits to view or interact with ads or generate traffic on ad-supported sites.

Inclusion Lists: identifies only the sites on which an advertiser wants its advertising to appear.

Interactive Advertising Bureau (IAB): An organization that develops industry standards, conducts research, and provides legal support for the online advertising industry.

Invalid Traffic (IVT): Online traffic generated from machines or other bot activity that interacts with digital ads.

Location fraud: Process in which fraudsters send false location information so the ads are served in locations other than where initially intended

Machine Learning: Scientific study of algorithms and statistical models that computer systems use to perform a specific task without using explicit instructions, relying on patterns and inference instead. It is seen as a subset of artificial intelligence.

Malicious apps: Apps that generate fraudulent impressions without the user knowing

Media Quality Report: Report released by IAS for global viewability, time-in-view, ad fraud, and brand safety benchmarks.

Open Measurement SDK (OM SDK): Offers common code and libraries for facilitating third-party access to measurement data

Pixel stuffing: The process of serving one or multiple ads in a single 1X1 pixel frame so that the ad can't be seen.

Proxy traffic: Traffic is routed through an intermediary proxy device or network where the ad is rendered in a user's device where there is a real human user.

Real Time Bidding (RTB): A means by which advertising inventory is bought and sold on an impression basis, via instantaneous auction, similar to financial markets.

Retargeting fraud: Bots mimic a human's intentions, such as an interest in a specific brand, in order to warrant the higher CPM typically associated with retargeting. Advertisers are deceived into believing they are receiving valuable, targeted audiences.

SDK spoofing: The act of hacking software development kits (SDK) and manipulating ad impressions by mimicking real user behavior.

Software Development Kits (SDK): A collection of software development tools in one installable package.

Sophisticated bot: A bot not listed in the industry bots and spider list or the known browser list.

Sophisticated Invalid Traffic (SIVT): Nonhuman traffic that is more difficult to detect and requires advanced analytics, multipoint corroboration/coordination, or significant human intervention to analyze and identify. Key examples include hijacked devices, hijacked tags, adware, malware, incentivized browsing, misappropriated content (if applicable), falsified viewable impression decisions, and cookie stuffing.

Traffic sourcing/Sourced traffic: Any method by which publishers acquire more visitors through third parties.

Trustworthy Accountability Group (TAG): An advertising industry initiative to fight criminal activity in the digital advertising supply chain. Through a cross-industry joint initiative, the IAB, the 4A's, and the ANA formed TAG to combat malware, fight Internet piracy, eliminate fraudulent traffic, and promote transparency.

User-agent spoofing: When a webpage header is modified to obfuscate information about the browser being used, which can interfere with user targeting.

Volunteer botnets: a distributed network of computers whose processing power is harnessed to steal data or carry out a political or socially-motivated denial of service (DoS) attack.

About IAS

Integral Ad Science (IAS) is the global market leader in digital ad verification, offering technologies that drive high-quality advertising media. IAS equips advertisers and publishers with both the insight and technology to protect their advertising investments from fraud and unsafe environments, as well as to capture consumer attention, and drive business outcomes. Founded in 2009, IAS is headquartered in New York with global operations in 18 offices across 13 countries. IAS is part of the Vista Equity Partners portfolio of software companies. For more on how IAS is powering great impressions for top publishers and advertisers around the world, visit **integralads.com**.

Integral Ad Science