Ad fraud

the essentials

IAS Integral Ad Science

Table of contents

INTRODUCTION

The IAB and Ernst & Young reported in 2015 that invalid traffic cost the U.S. digital marketing, advertising, and media industry around $4.6 billion annually.

For more findings, check out the full IAB report.

In our year-end survey, we surveyed members of the digital advertising industry—agencies, brands, DSPs, networks, publishers, and trading desks—to better understand their concerns for the year ahead.

Not surprisingly, ad fraud topped the list.

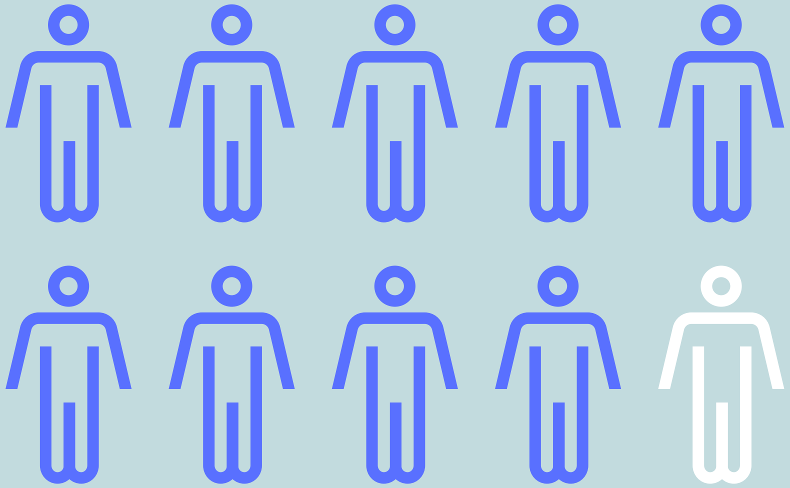

89%

said it has a direct impact on media quality

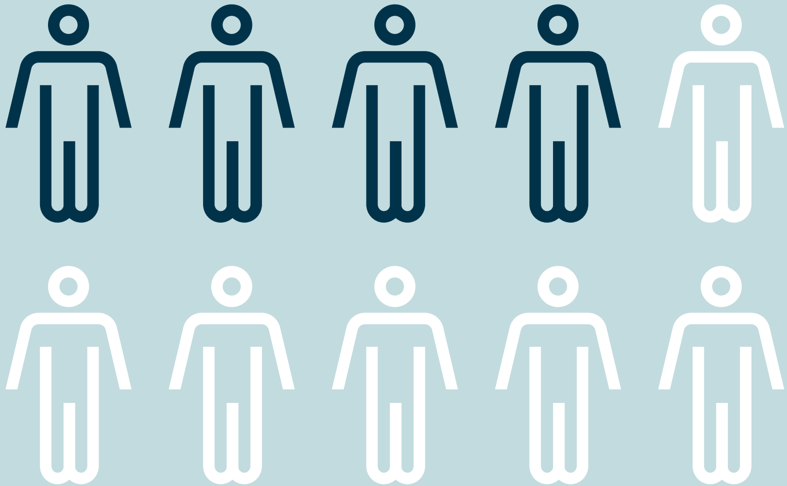

39%

think it trumps other ad-quality measurement factors like brand safety, viewability, transparency, and geo-compliance

And yet, only 43% said they understand how fraud is detected.

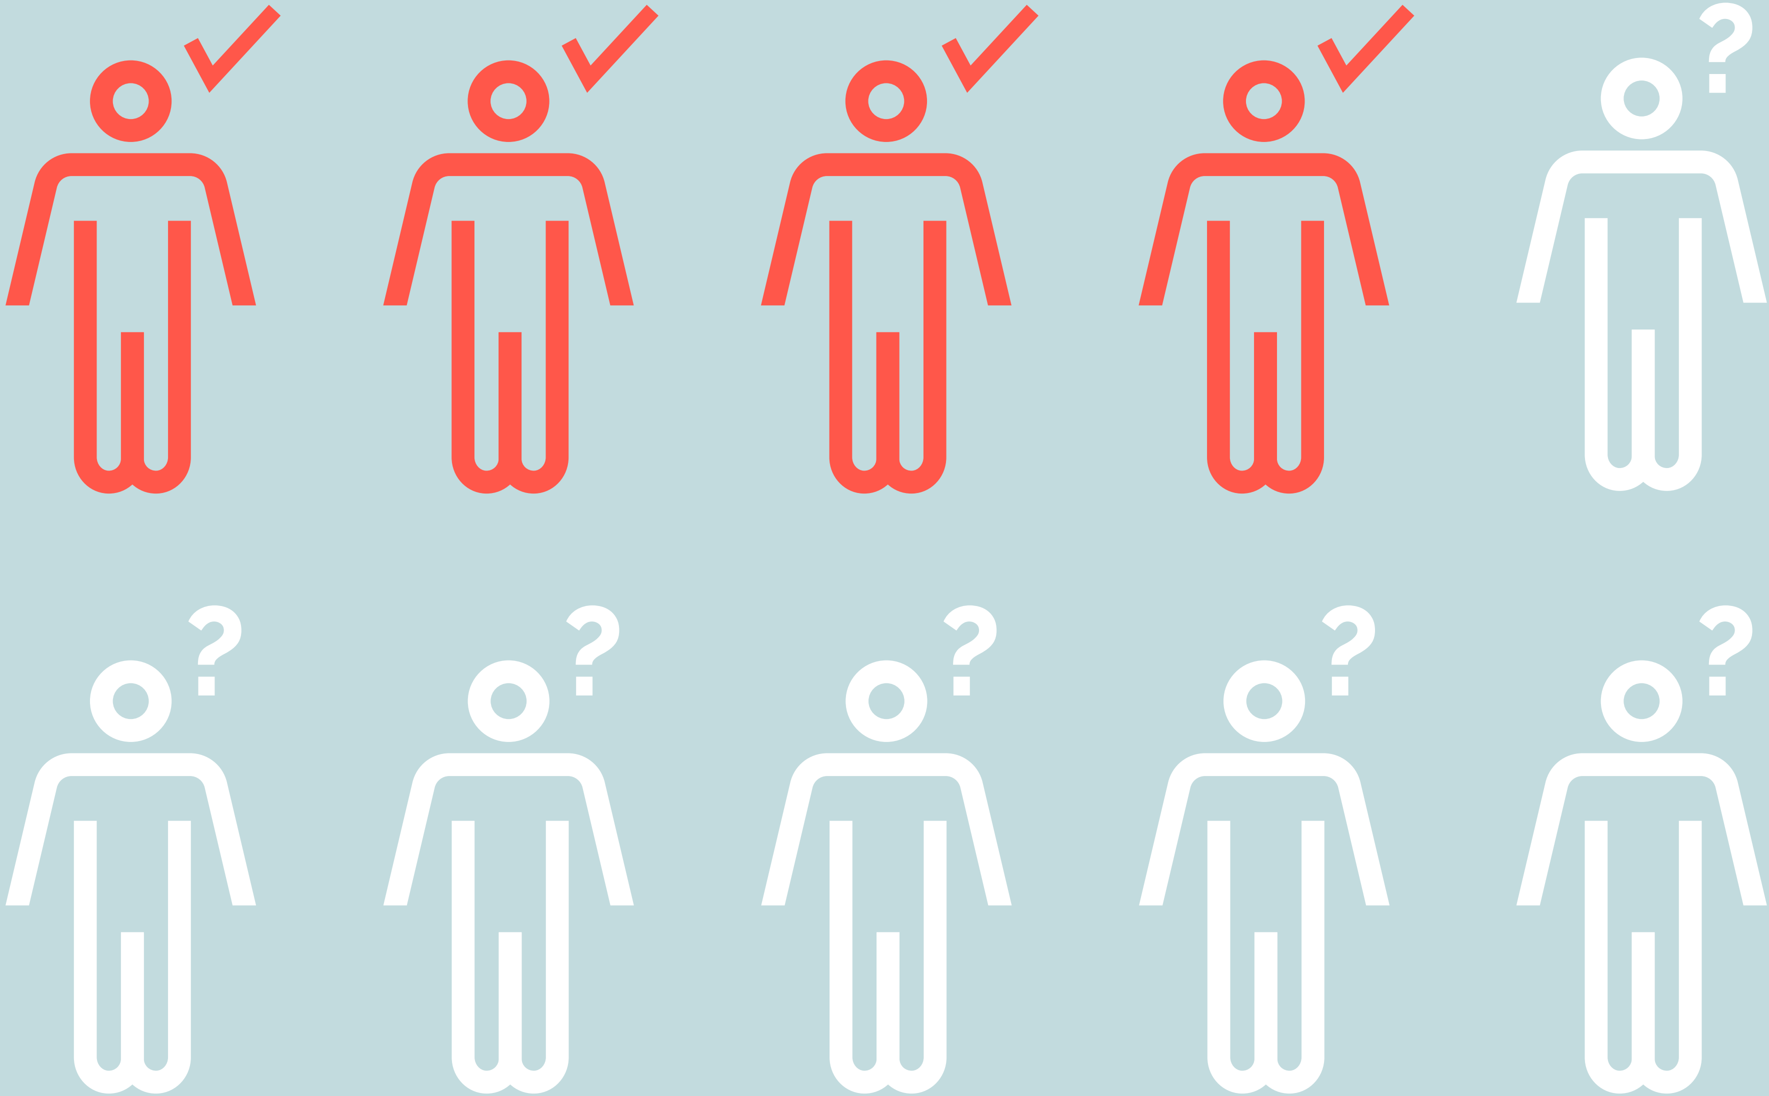

Everyone knows there is a problem, but it’s not always clear:

What exactly is fraud

How it works

How to eliminate it from the digital ecosystem

We’re here to make this complex topic easy to understand and address.

CHAPTER ONE:

Overview

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.

Most often, fraud refers to certain kinds of traffic, not to publishers or ad tech partners that are a part of the supply chain. There are publishers with high proportions of fraudulent traffic, and others with very low proportions.

The landscape of fraud is ever-changing: fraud may concentrate on one site one week and somewhere else the next. Even premium publishers can be subject to hit-and-run attacks. Every traffic source requires constant re-evaluation.

So, what *exactly* is ad fraud?

Selling inventory automatically generated by bots or background mobile-app services

Serving ads on a site other than the one provided in an RTB request—this is known as domain spoofing

Delivering pre-roll video placements in display banner slots

Falsifying user characteristics like location and browser type

Hiding ads behind or inside other page elements so that they can't be viewed

Hindering a user's opportunity to engage by frequently refreshing the ad unit or page

The most prevalent forms of fraud are:

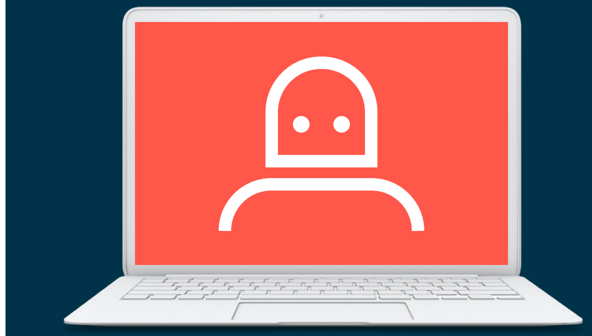

Bots

The most common—and well-known—example of fraud is bots. Because bots are such a large and prevalent issue, we've devoted an **entire section** to it.

www.adsafe.org
BID
ACTUAL WEBSITE
www.safeads.com

Domain spoofing

Occurs in a real-time bidding (RTB) environment, where the URL is used to fool an agency into thinking their ad is going to a premium site, when instead it's going to a low-quality website—or that their ad is going to a brand-safe site when it's actually going to a brand-unsafe site. To learn more about domain spoofing, check out the **types of ad fraud** section below.

Why?

These are 0-to-60 forms of fraud. With relatively little effort or expertise, traffic that previously would have been worth nothing at all can be sold at high CPMs. Other forms of fraud are either technically challenging or provide smaller boosts to CPMs.

What *isn't* ad fraud?

Self-declared web crawlers and other good bots

Collisions (ads from the same brand accidentally appearing on the same page)

Poor viewability

Muted video

How pervasive is ad fraud?

Out of the nonhuman traffic that cost the industry $4.6 billion annually according to the IAB report:

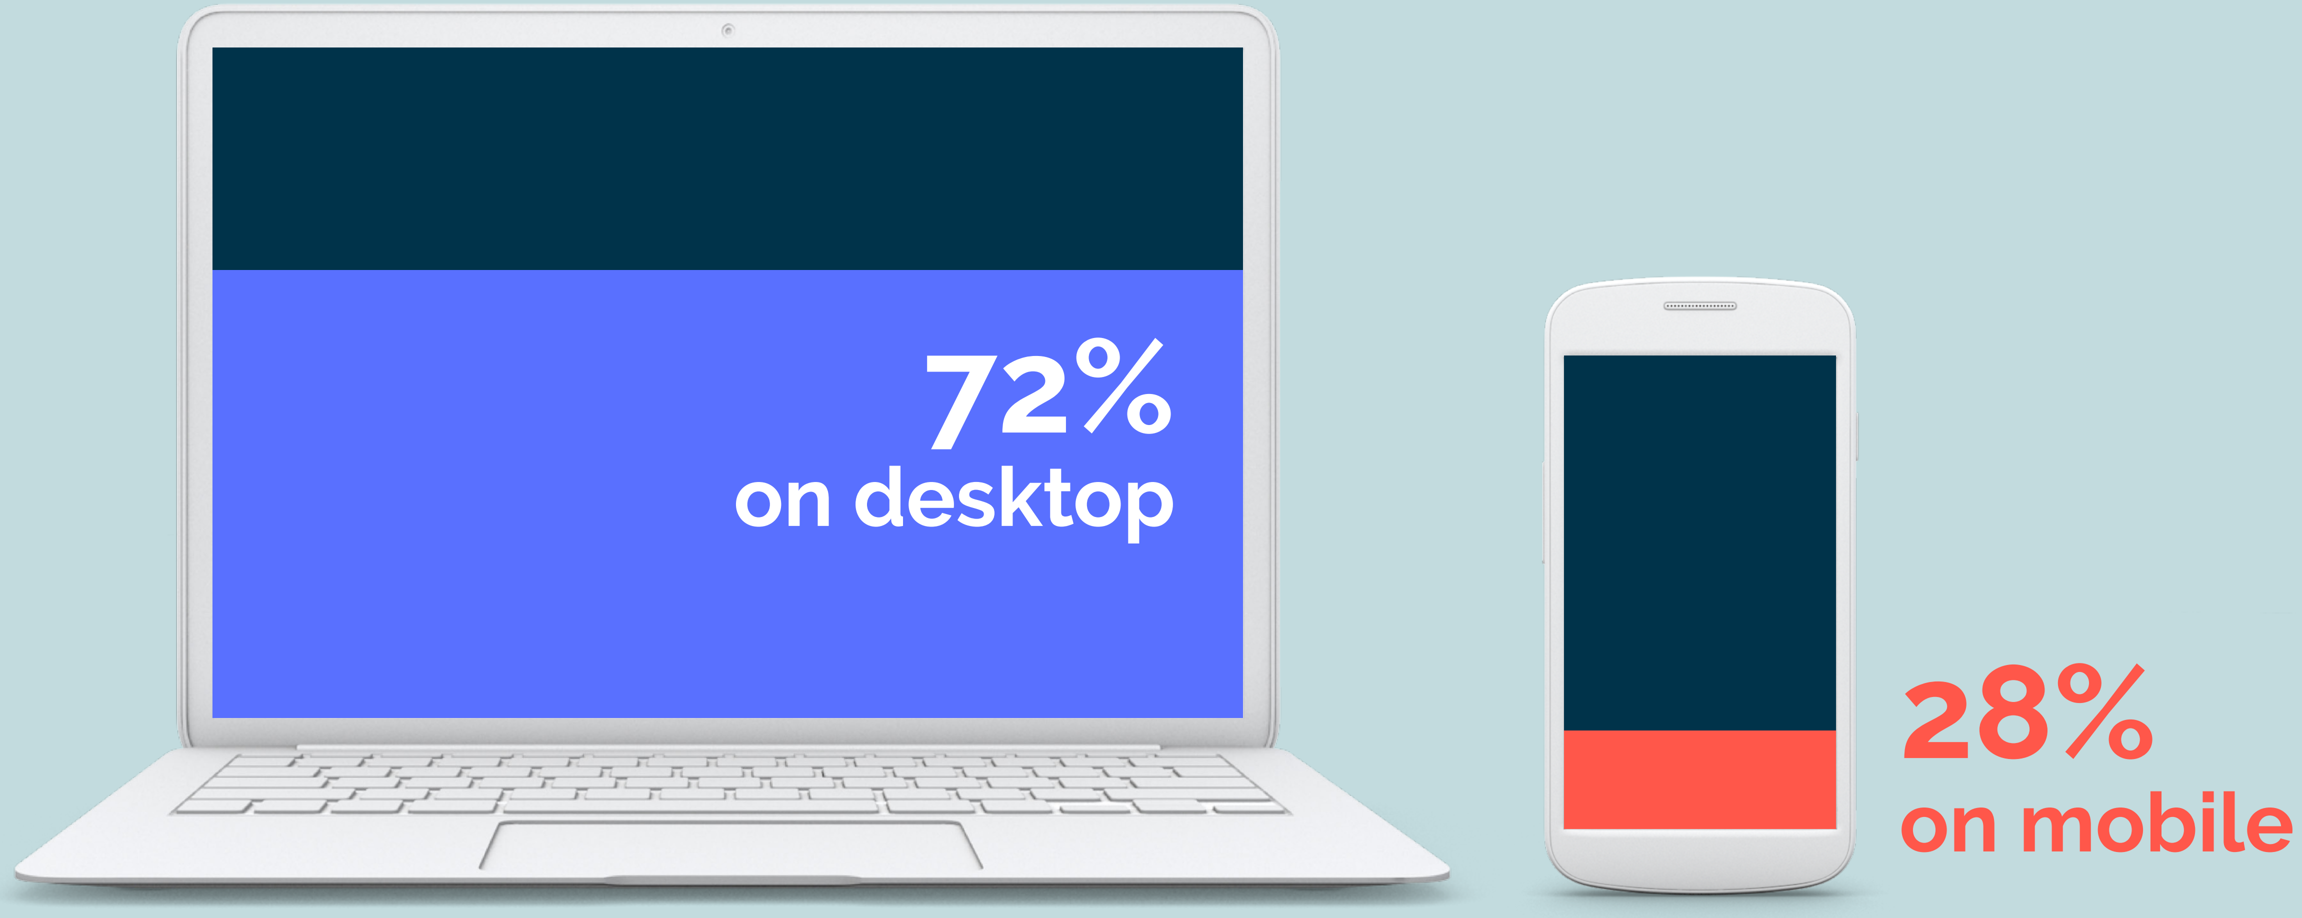

Inventory types that often have higher rates of fraud include:
Programmatic & **Video**

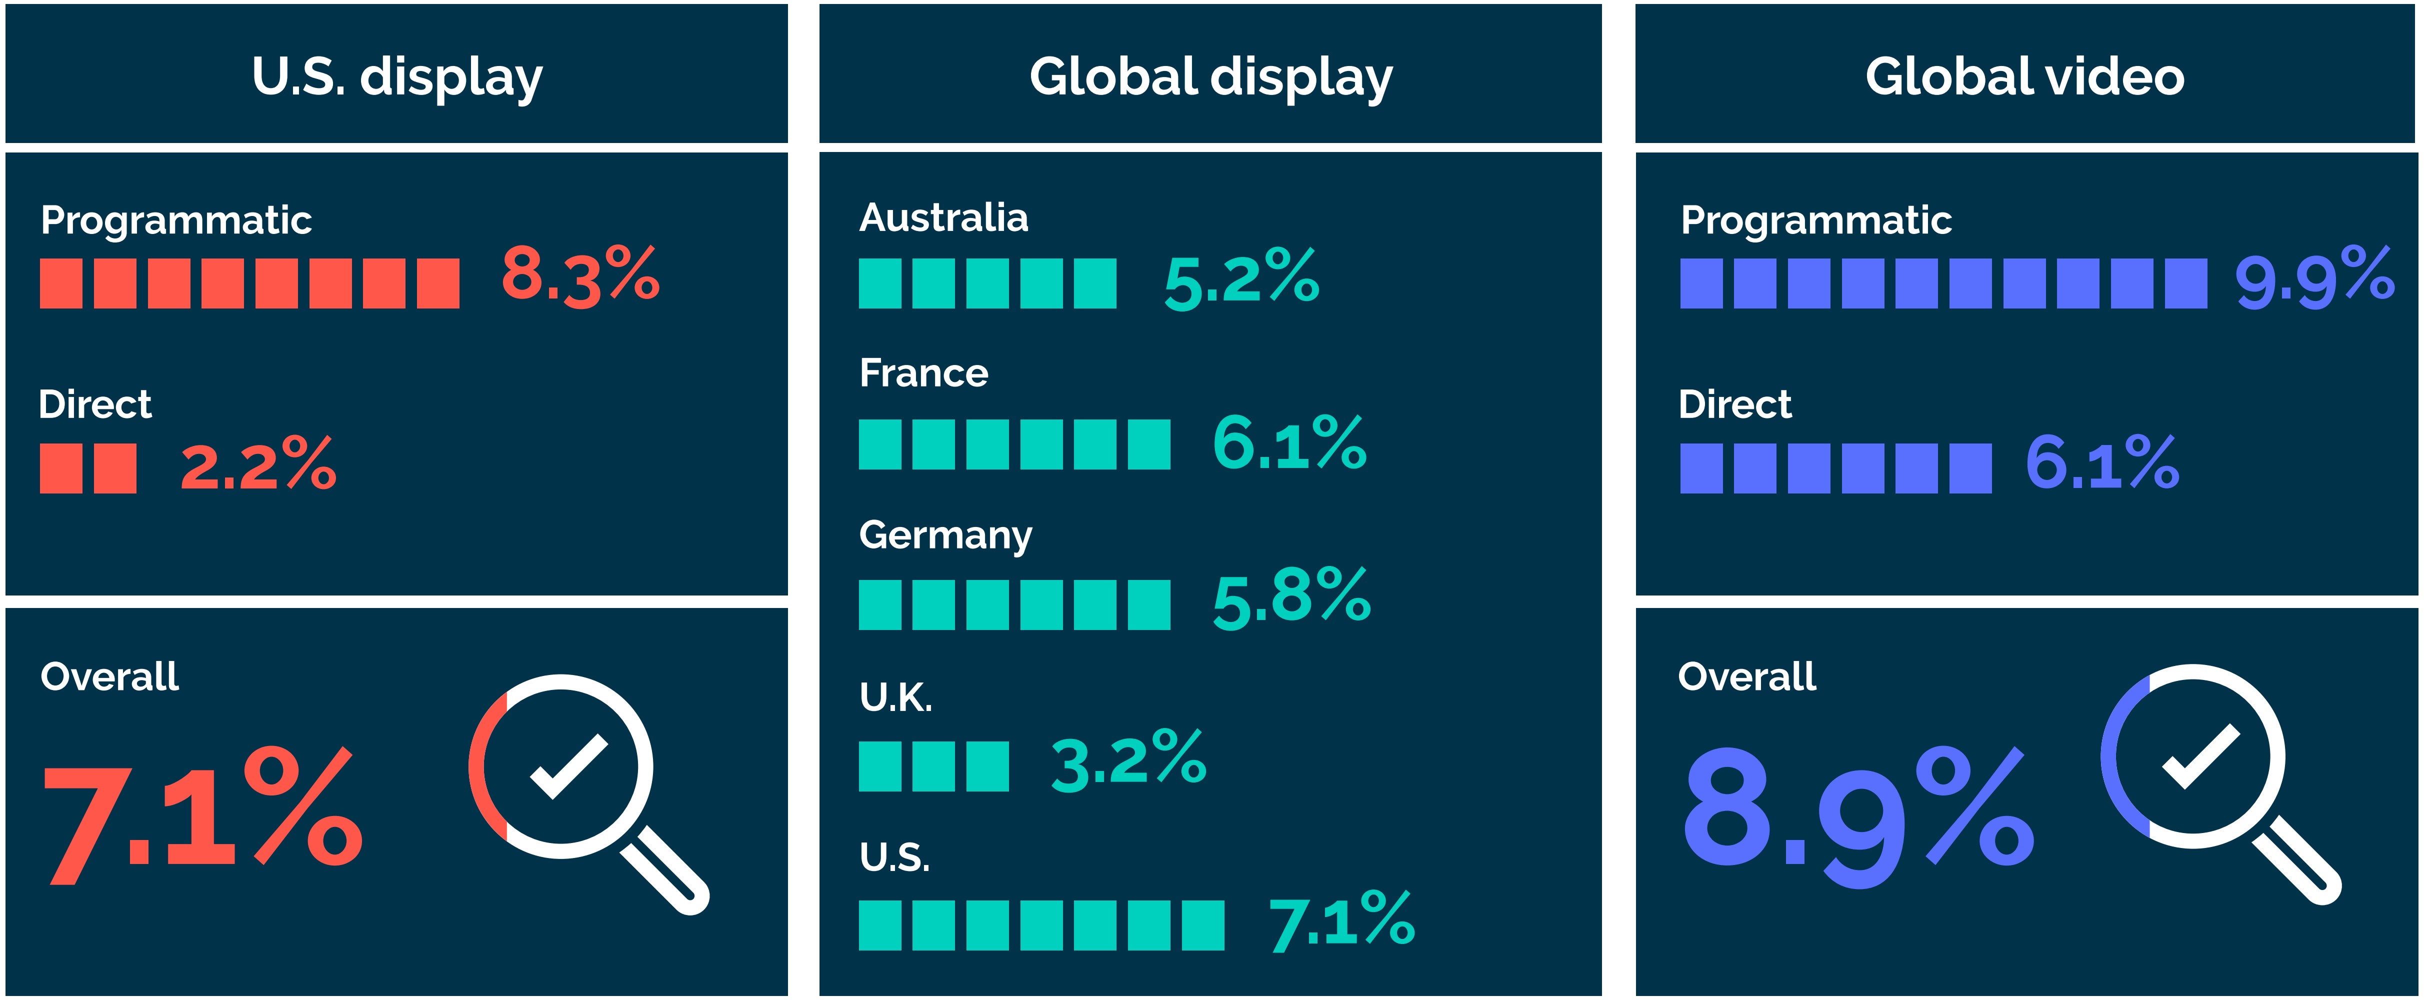

Across desktop display, 2-4% fraud is typical for direct buys, while programmatic buys can reach 4x as much, usually 10-15%. Of course, if people aren't using any sort of fraud detection or prevention, rates can balloon past 50% on both direct and programmatic.

For more highlights on the state of fraud in digital advertising, check out our H1 2016 Media Quality Report

Why does ad fraud occur?

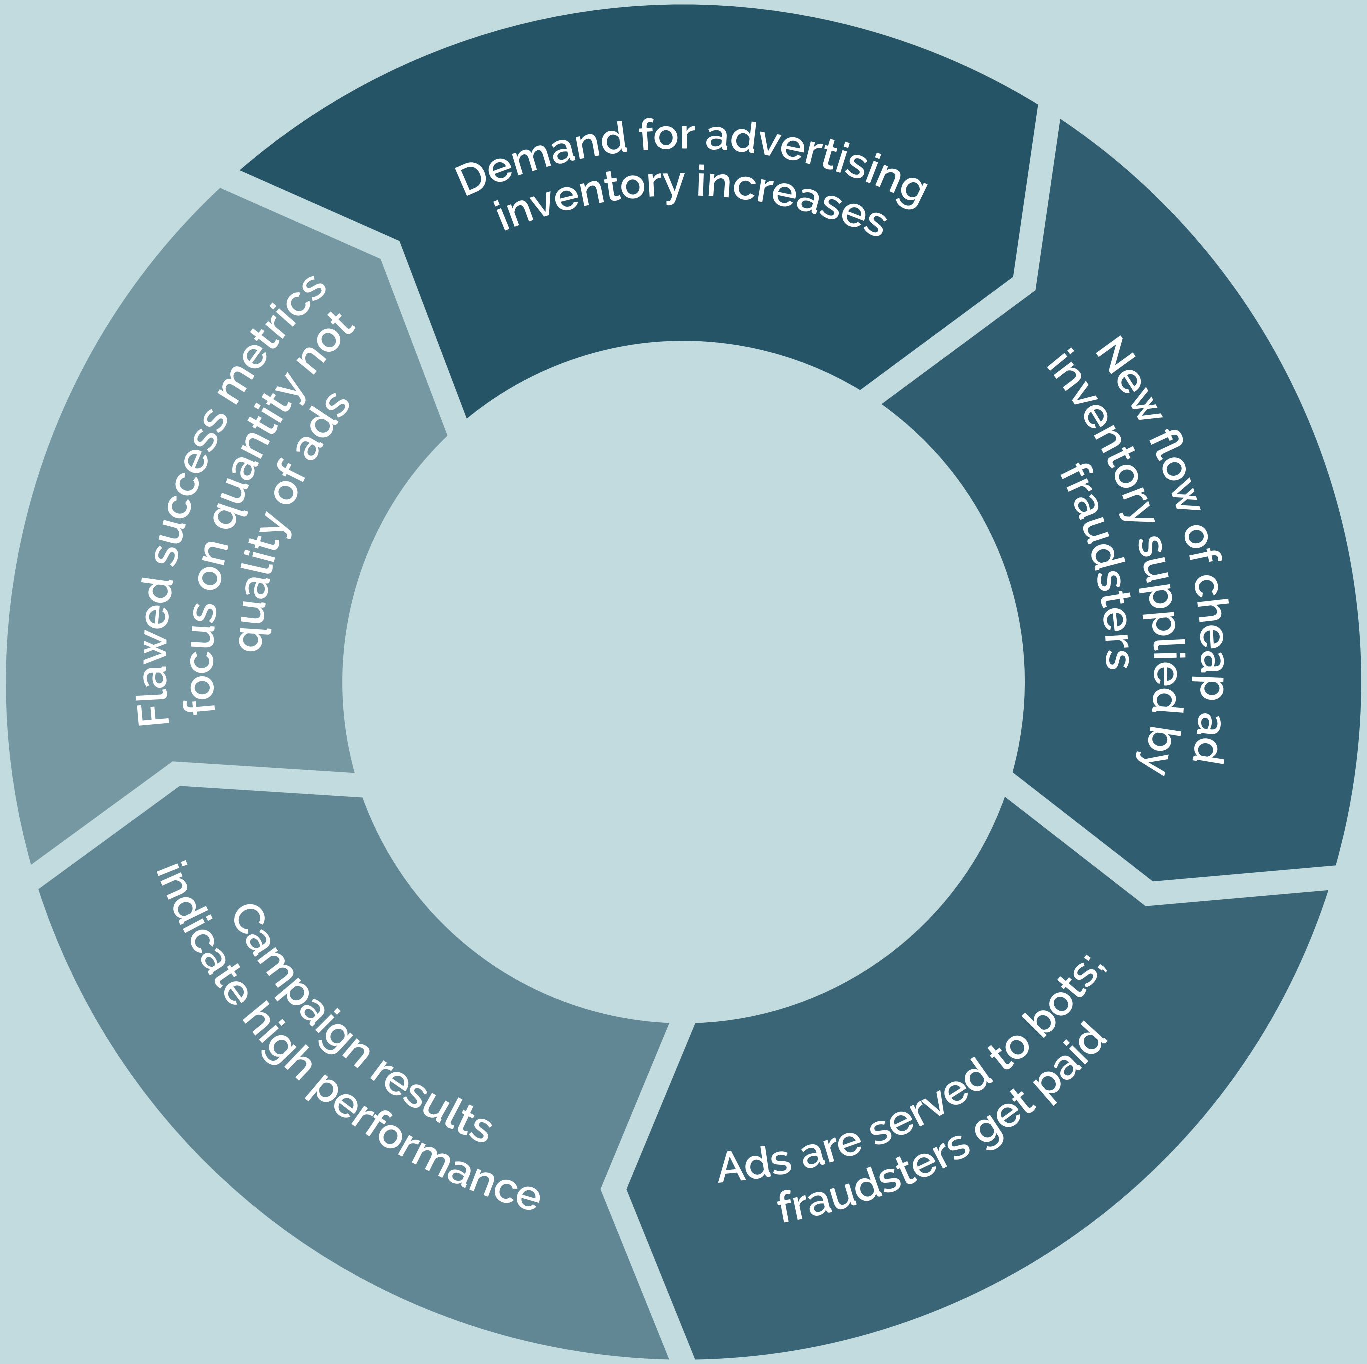

Demand for advertising inventory increases

New flow of cheap ad inventory supplied by fraudsters

Ads are served to bots; fraudsters get paid

Campaign results indicate high performance

Flawed success metrics focus on quantity not quality of ads

How does fraudulent traffic occur?

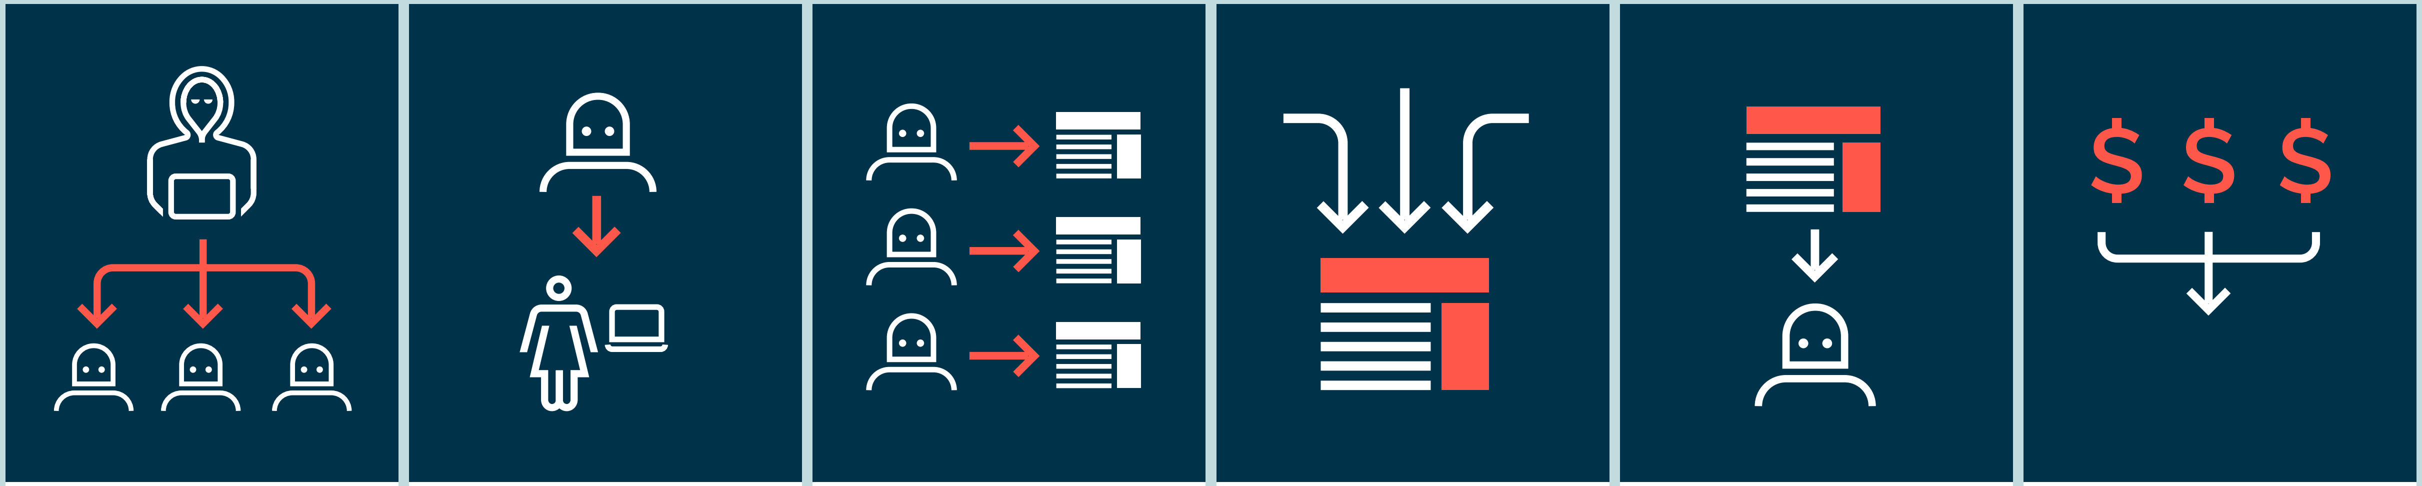

Hackers use code to create bots able to take orders from botnet center

Users unknowingly download and install bot engines 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

Who's involved?

All parties that go through the process of delivering a given fraudulent impression—publisher, network, exchange, traffic broker, malware distributor—are so interconnected that it's nearly impossible to determine who's at fault.

The expected parties behind fraud are hackers and botnet operators.

Hacker

- Sex: Male
- Age: 18-35
- Location: Eastern Europe, Asia
- Background: Good computer skills

Botnet operator

- Sex: Male
- Age: 34+
- Location: Eastern Europe
- Characteristics: Disregard of the law, confident, driven by money

Typically infected computer owner

- Average technology skills
- May own a dated computer or software

Fraudsters often operate out of areas where it's tough to enforce ad fraud laws, such as Eastern Europe, Russia, and Asia.

Hackers can be motivated by money, curiosity, notoriety—but the botnet operator is in it for the money. These operators orchestrate the ad fraud using hackers to their advantage. Hackers write the code to break into a computer and take control of it. The botnet operator coordinates with traffic brokers and issues instructions for the bots to follow.

Hackers and botnet operators could be the same person, but often are not.

Another key player is the traffic broker. These middlemen connect websites looking to boost their traffic to the botnet operators who can supply it. They also frequently sell to each other in a complicated web of arbitrage.

An unfortunate reality is that there are many people with ties to the ad tech industry who are familiar with fraud and how to use it to make money for themselves, from establishing bogus content networks to funneling trade secrets to hackers to using a legitimate business as a front for traffic selling. A key challenge in rooting out fraud is ensuring cooperation throughout the ecosystem.

For more on the basics of fraud, watch this presentation by David Hahn, our E.V.P. of Strategy.

Key takeaways

Programmatic buys often attract more fraud than direct buys.

Video ads often also have greater instances of fraud than other digital formats.

For more on video and fraud, check out the section below.

Wherever ad spend is growing at a rapid pace—like digital video—you'll see more demand than supply. That's where fraudsters hit.

What are the industry guidelines?

Media Rating Council (MRC) guidelines

The MRC guidelines on **Invalid Traffic (IVT)** were published in October 2015, to address inconsistent methods of removing invalid traffic and the resulting irreconcilable discrepancies. They establish minimum requirements for identifying and removing invalid traffic from advertising transactions.

IVT induces systems to generate actions that take away from the proper delivery of ads to the right people at the right time. It can impact display, video, mobile, audio, search, and social. It can include forms of legitimate activity as well as activity generated by bad actors for malicious purposes.

The MRC guidelines created two categories: **general** and **sophisticated** invalid traffic.

General IVT

General IVT consists of traffic that can be identified through routine means of filtration, executed by using lists or other standardized parameter checks.

Examples:

- Traffic from datacenters (this traffic is usually nonhuman)
- Spiders and other crawlers pretending to be legitimate users
- Bots detected through simple activity-based metrics like impossibly high impression volumes

Sophisticated IVT

Sophisticated IVT consists of traffic that is more difficult to detect, requiring advanced analytics, multipoint corroboration/coordination, significant human intervention, etc., to analyze and identify.

Examples:

- Falsely represented sites or impressions
- Hijacked devices: a user's device (browser, phone, app) is modified to request HTML or make ad requests that are not under the control of a user and made without the user's consent (for example, operations made by a bot)
- Hijacked sessions within hijacked devices
- Hidden/stacked/covered or otherwise intentionally obfuscated ad serving
- Anonymized proxy traffic
- Incentivized manipulation of measurements such as payment for video interaction or guided browsing
- Misappropriated content
- Falsified viewability measurement
- Cookie stuffing
- Manipulation or falsification of location data or related attributes

Cases of unintended IVT

Sometimes fraudulent invalid traffic can occur without the original parties meaning for it to happen. Some examples include:

Traffic sourcing

When publishers sell more inventory than currently available, or otherwise need to increase traffic to their site to meet advertiser needs, they may seek out third-party publishers or traffic-broker sites to purchase that additional traffic. However, these can be operated by fraudsters, who will deploy bots to drive up the numbers for the original publisher.

- Sourced traffic can be delivered in several ways, including driving bots to the site, hiding the site within another, and a technique for domain spoofing in which one site loads another's ad units.
- This approach has very low visibility, so it's extremely difficult to know who was behind the fraud, and which traffic was real and which was bought. And traffic brokers frequently sell to each other in a web of arbitrage, compounding the difficulty.

Audience extension

When publishers fulfill their ad buys with inventory placed on other sites they own, there can be a lack of transparency which can lead to serving ads outside of the target audience.

Cost of IVT

Lowered inventory CPMs

Damaged reputation of organizations susceptible to fraudulent IVT

Reluctance to invest and allocate digital media spend

Monetary costs to fight it

Key takeaways

1

It’s hard to identify and prevent its monetization

2

It’s difficult to maintain transparency into the quality of audience fulfilling the buys

3

The rise of automation and increased complexity in the digital workflow means the prevalence of IVT will persist

4

Publishers purchasing low-cost traffic on open ad exchanges can lead to more fraudulent IVT in transactions

5

Has direct monetary impact on the buy side if they don’t protect their inventory and reputation

CHAPTER TWO:

Types of fraud

Ad fraud disrupts the aim of advertising: delivering the right *message*, to the right person, in the right *place*.

Fraudsters compromise all three areas of advertising through various techniques like pixel stuffing, ad stacking, nonhuman traffic, domain spoofing, user-agent spoofing, and more.

The most prevalent forms of fraud are nonhuman traffic and domain spoofing.

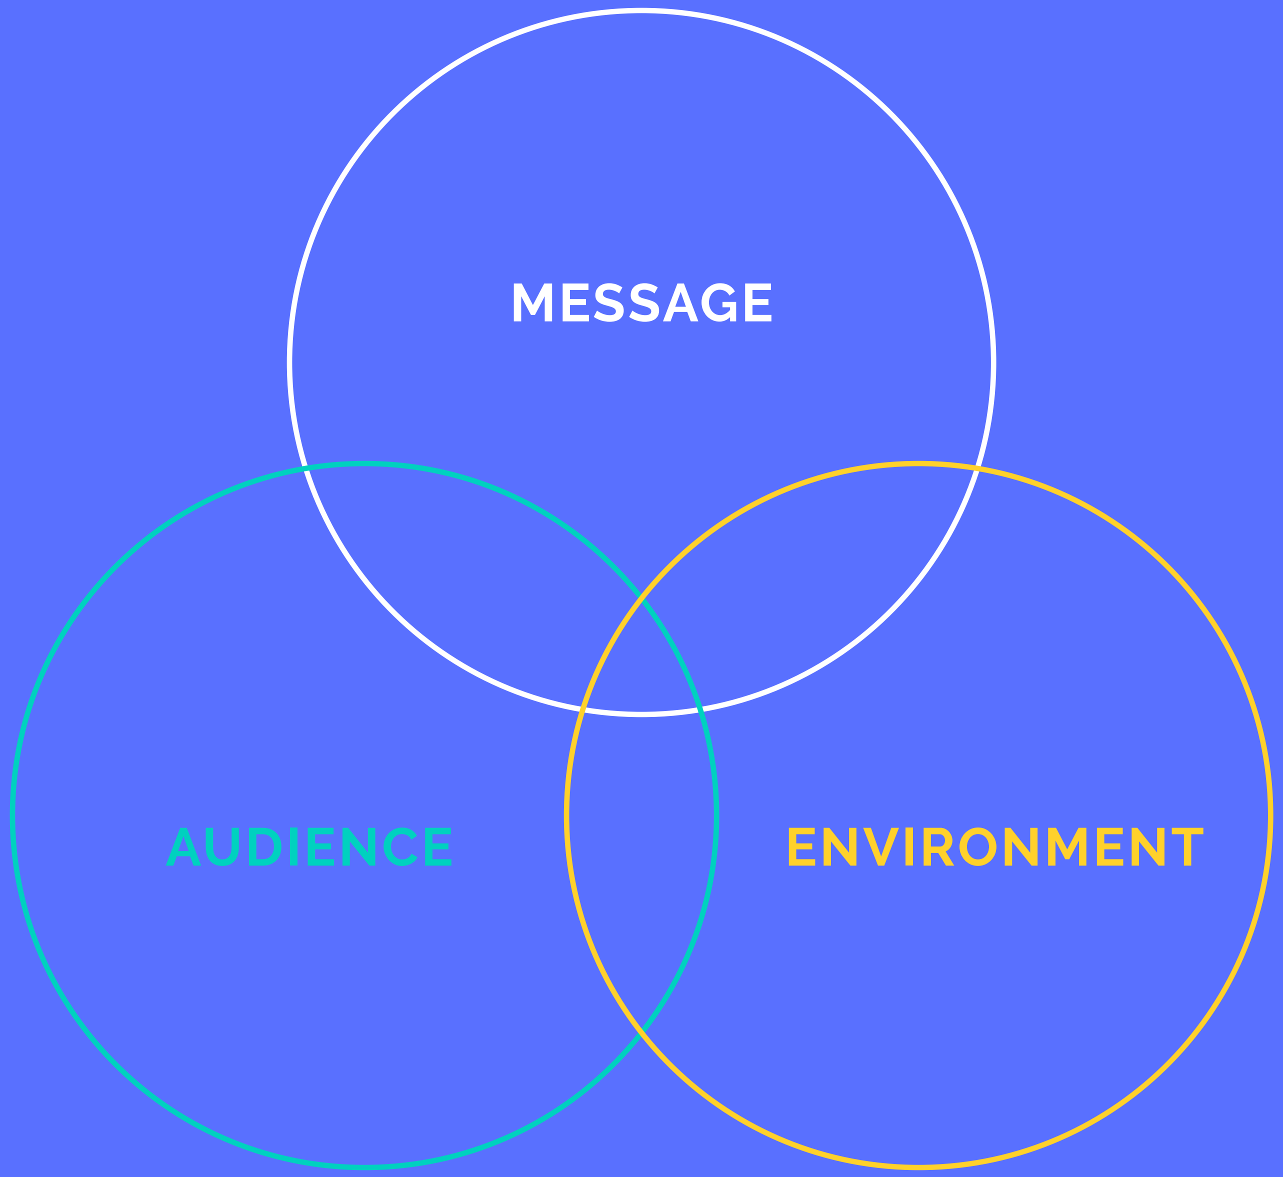

Pixel stuffing

Serving one or more ads or an entire ad-supported site in a single 1x1 pixel frame, so that the 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. (Imagine a stack of pancakes.) The advertiser is paying for impressions even if the end user is not seeing an ad.

Location fraud

Location is a critical part of media plans for advertisers. Since agencies are willing to pay a premium for location data (in order to geo-target appropriately), fraudsters will send false location information.

In location fraud, an advertiser pays a premium CPM for inventory to be served in a particular country or region, but the traffic is actually served elsewhere. For users, they might be surfing the web on their mobile device in NYC, and all of a sudden see ads for Alcatraz – which would of course yield few conversions.

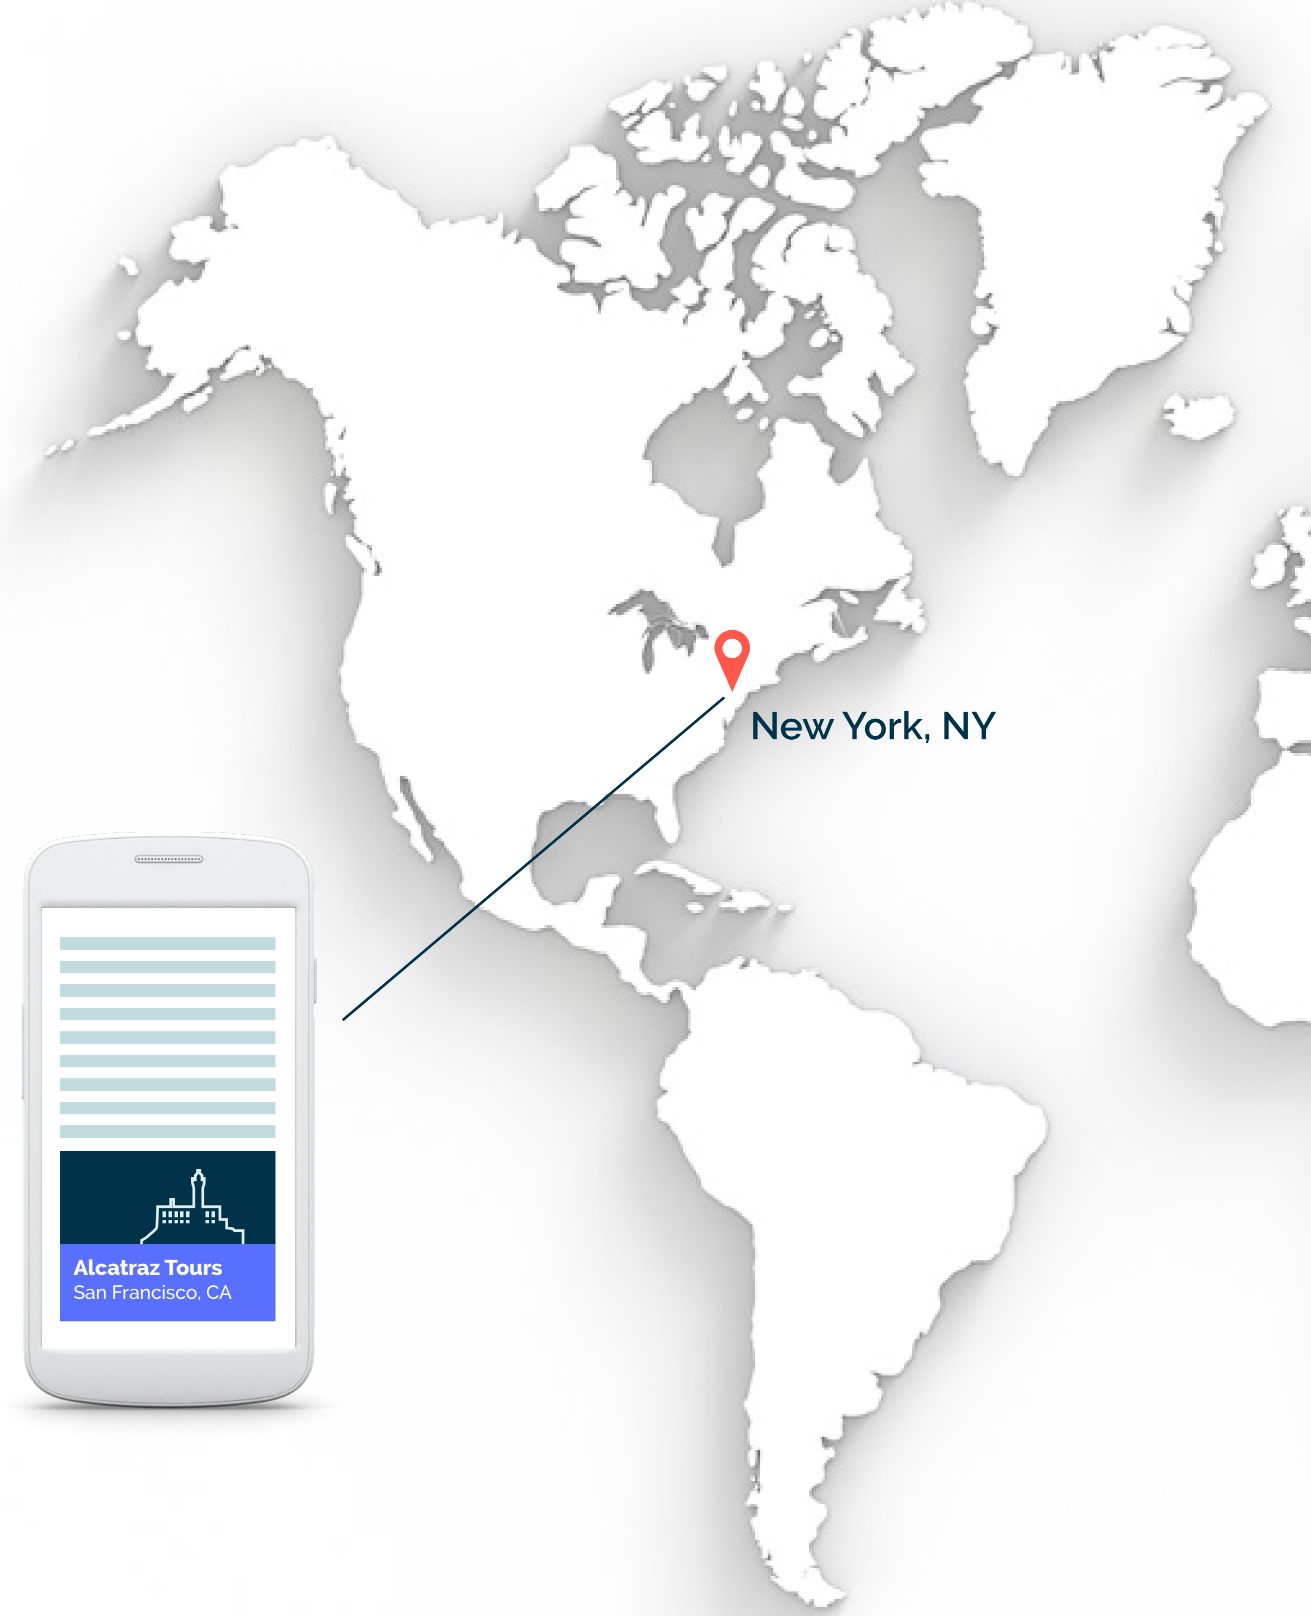

Cookie stuffing

Cookies are a vital method of tracking user behavior, and ultimately help determine what advertising effort led to a conversion (click-through, 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 down the line, the website associated with the stuffed cookie gets credit – and profit – for that action. This could also happen on a broader scale with a network of sites placing each other's cookies on users and then sharing the pot later.

Cookie stuffing can also refer to the practice of placing many cookies on a user or bot so that they get targeted at higher CPMs, even if they haven't visited sites that indicate they are potential high-value consumers.

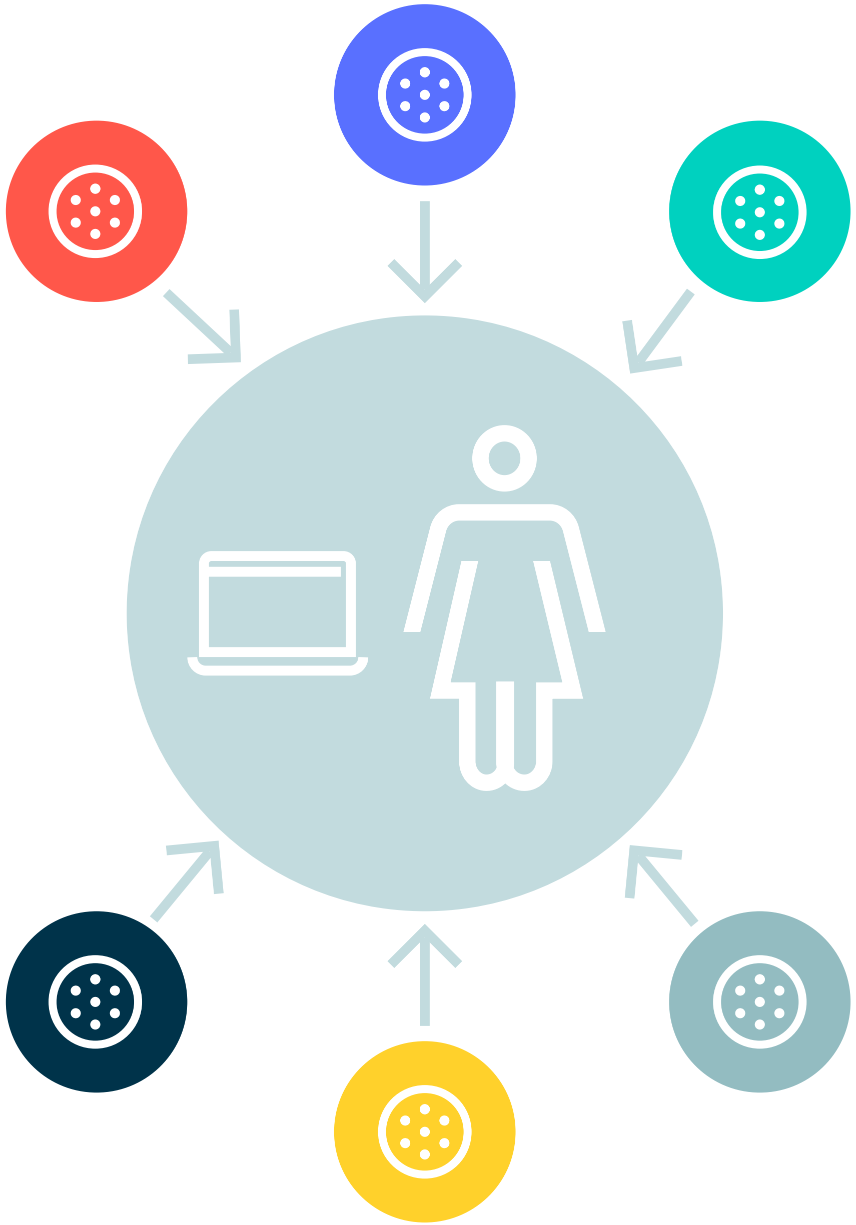

User-agent spoofing

Every request for a web page is sent with a “header” that provides some basic information about it, including where it came from, what language it is expecting, the time, and other pieces of information. One piece of information is a detailed description of the browser: its type, version, operating system, even plug-ins.

In spoofing, the information is modified to lie about the browser that’s being used, which can interfere with some kinds of user targeting. It’s most often seen with bots trying to hide their tracks, but some human users will occasionally engage in this as well.

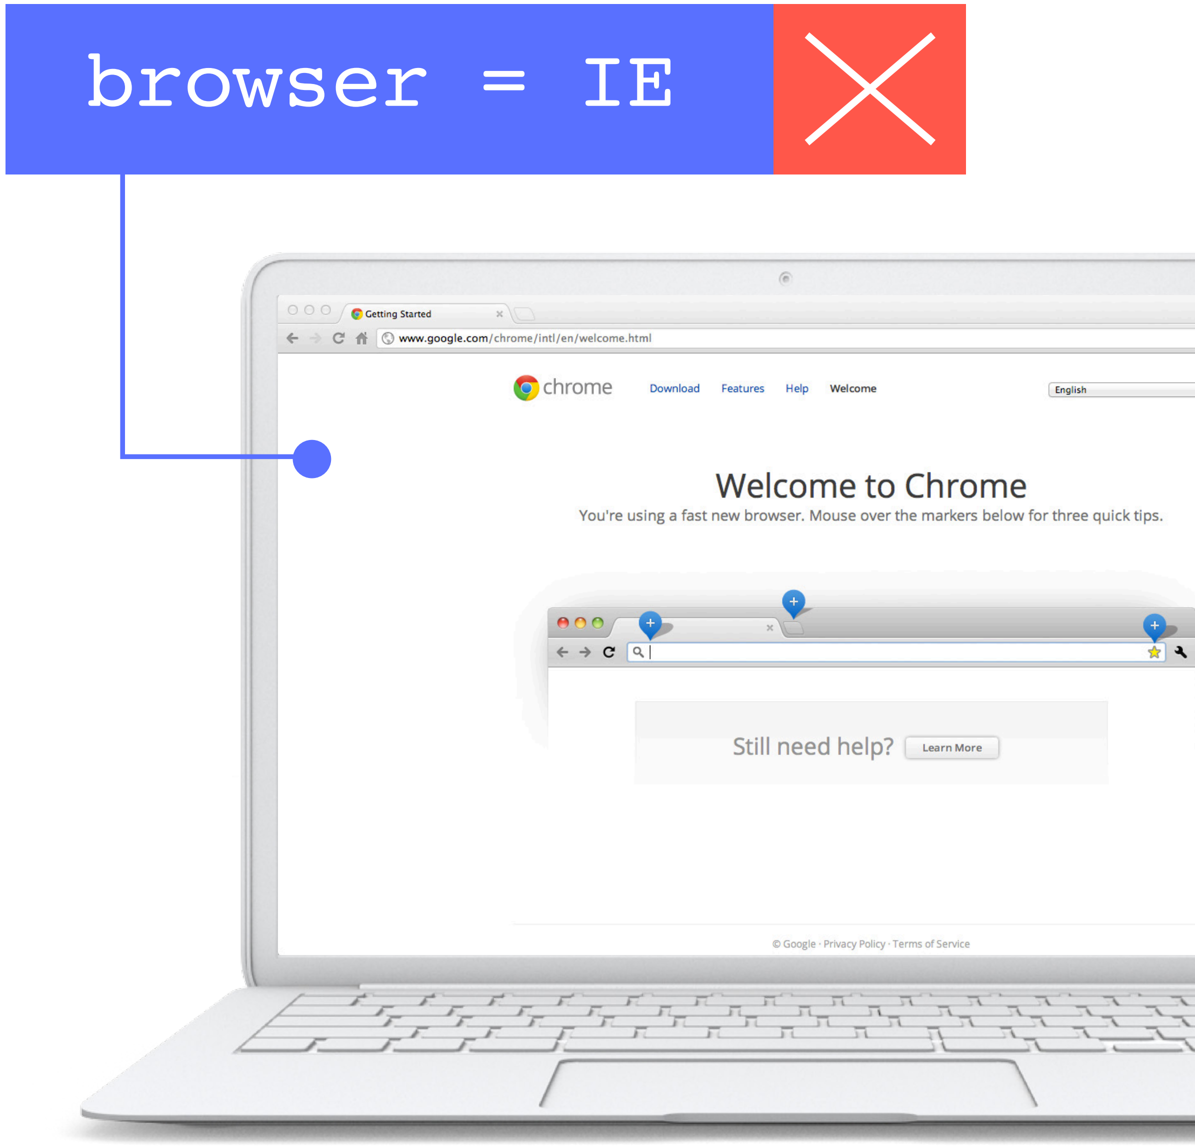
browser = IE
Getting Started
www.google.com/chrome/intl/en/welcome.html
chrome
Download
Features
Help
Welcome
English
Welcome to Chrome
You're using a fast new browser. Mouse over the markers below for three quick tips.
Still need help?
Learn More
© Google - Privacy Policy - Terms of Service

Domain spoofing

Occurs in a real-time bidding (RTB) environment, where the URL is used to fool an agency into thinking their ad is going to a premium site, when instead it's going to a low-quality website—or that their ad is going to a brand-safe site when it's actually going to a brand—unsafe site. The impressions and users are real, but the inventory is falsely represented, and therefore purchased at much higher CPM rates. This hurts both the buy and sell sides.

Domain spoofing is also commonly used to mask unsafe sites. Brand safety is a huge concern to advertisers, and fraudsters can take advantage by spoofing the domains of sites like video piracy sites, etc., in order to conceal their real identity and monetize the traffic. This specific type of domain spoofing is called **cross-domain embedding**.

Domain spoofing can occur partly due to the reliance on using whitelists for brand safety.

For more on domain spoofing, check out this article on AdExchanger.

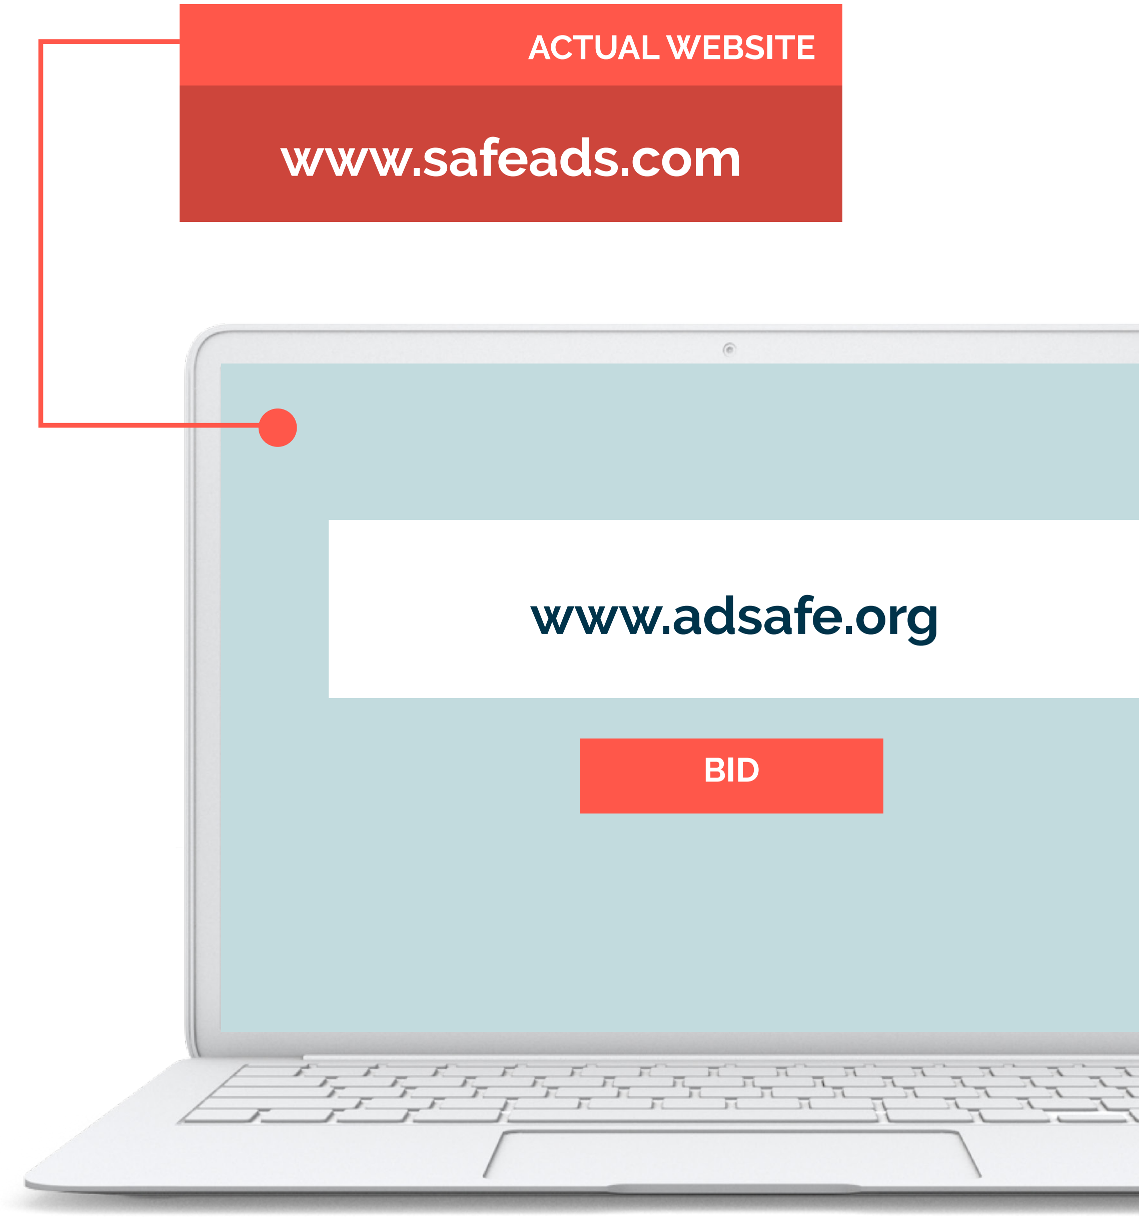

Costs of domain spoofing

- Loads programmatic buys with low-quality inventory
- Violates the security of whitelists
- Throws off KPIs
- Steals advertiser spend
- Steals publisher potential ad revenue

It’s important to note that 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.

Sites that look real

- Have content stolen from legitimate sites, often in verticals designed to attract higher CPMs: fashion, food, and news
- Receive invalid traffic from any source: bots, hidden ads, or cross-domain embedding
- Might themselves be filled with ads, which are hidden from your view but rack up even more impressions

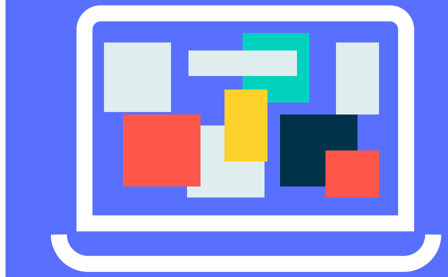

Sites that are only filled with ads

- Cause significant drain on device resources
- Are designed primarily for visits by bots, though they could be hidden in human browsers through pixel stuffing

To learn more about different types of fraud, check out Jason Shaw, our director of data science, talking about Ad Fraud: Beyond Bots.

CHAPTER THREE:

Bots

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

While other forms of fraud provide small boosts to CPMs, bot traffic can create revenue streams where there were none before. Bot traffic also makes it harder for the industry to identify who's behind it.

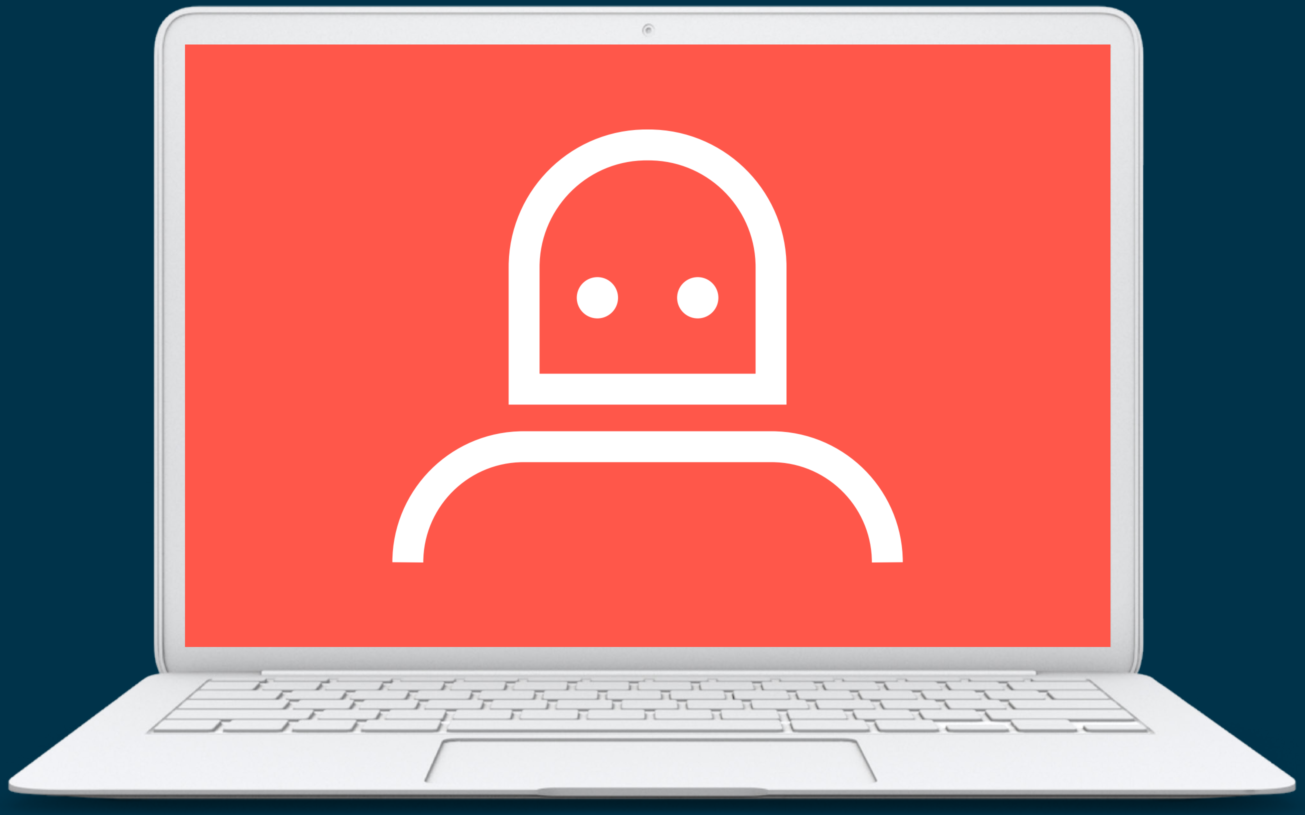

Hackers create bots to surf the web, click on ads, play videos, and perform searches, which all drive up traffic, resulting in more money for the fraudsters.

These bots are viruses that can be installed unknowingly on a computer and that use computer resources in a conservative manner so you wouldn't notice. For example, if you noticed it takes 30 seconds to load a page, you might suspect you have a virus. On the other hand, if it takes 10 seconds, you might think it's because your computer is just older and getting slower.

Most people with infected computers are completely unaware.

Bots vary in levels of sophistication and structure and perpetrate ad fraud in slightly different ways.

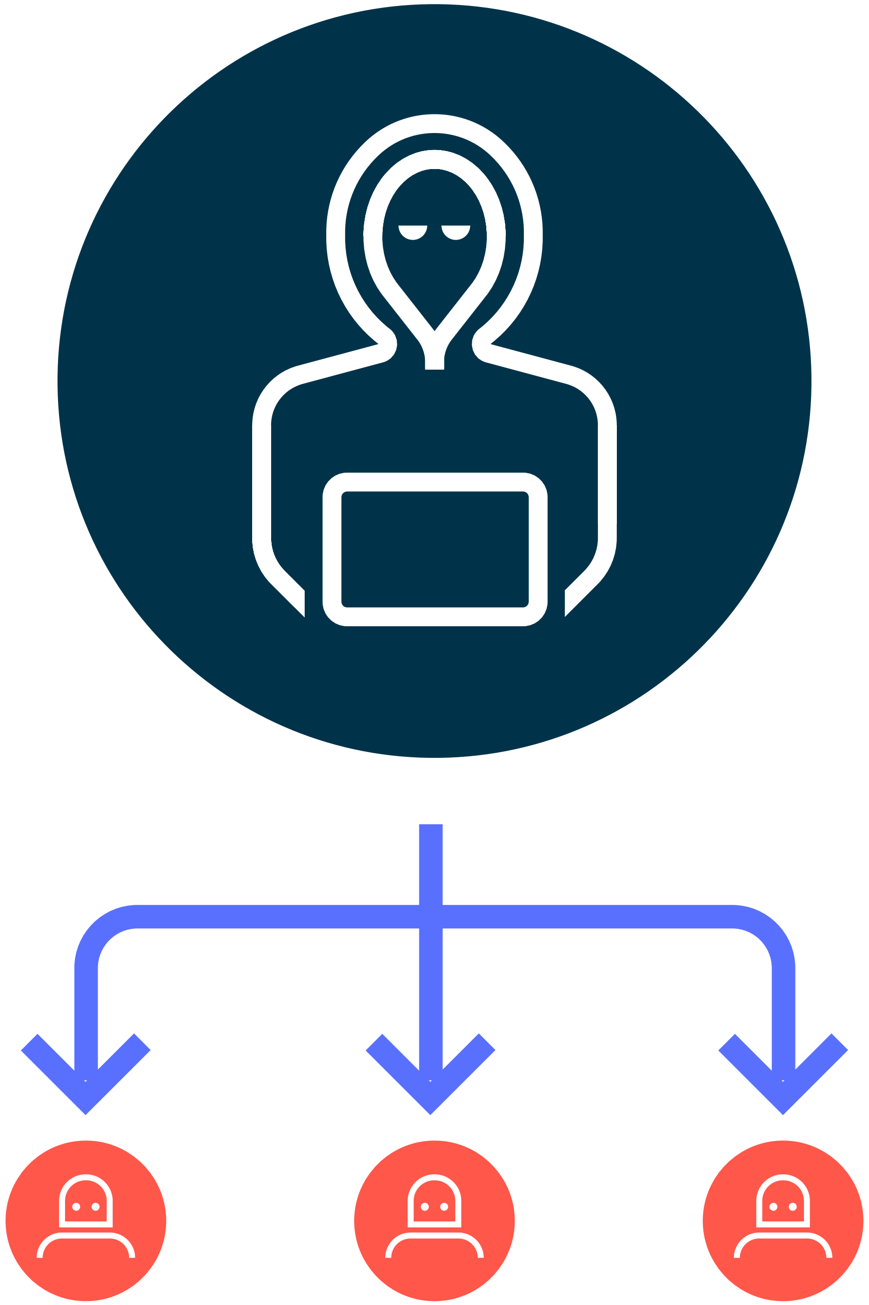

Fraudsters can create bot networks that perpetrate ad fraud with the compromised computers, all unbeknownst to the human user.

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

Most malware is built with the ability to join a botnet. The typical ways malware comes to infect a user's computer include: opening e-mail attachments, navigating to dangerous links, and installing software from untrustworthy sources.

Upon infection, the malware turns the computer into a bot, one part of a large network of infected machines. The bot begins communicating with a command and control (C&C) server, which gives instructions for the bot to follow. These can include activities like visiting various premium sites 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 and have attracted the same advertisers. Those ads are therefore wasted on the bots.

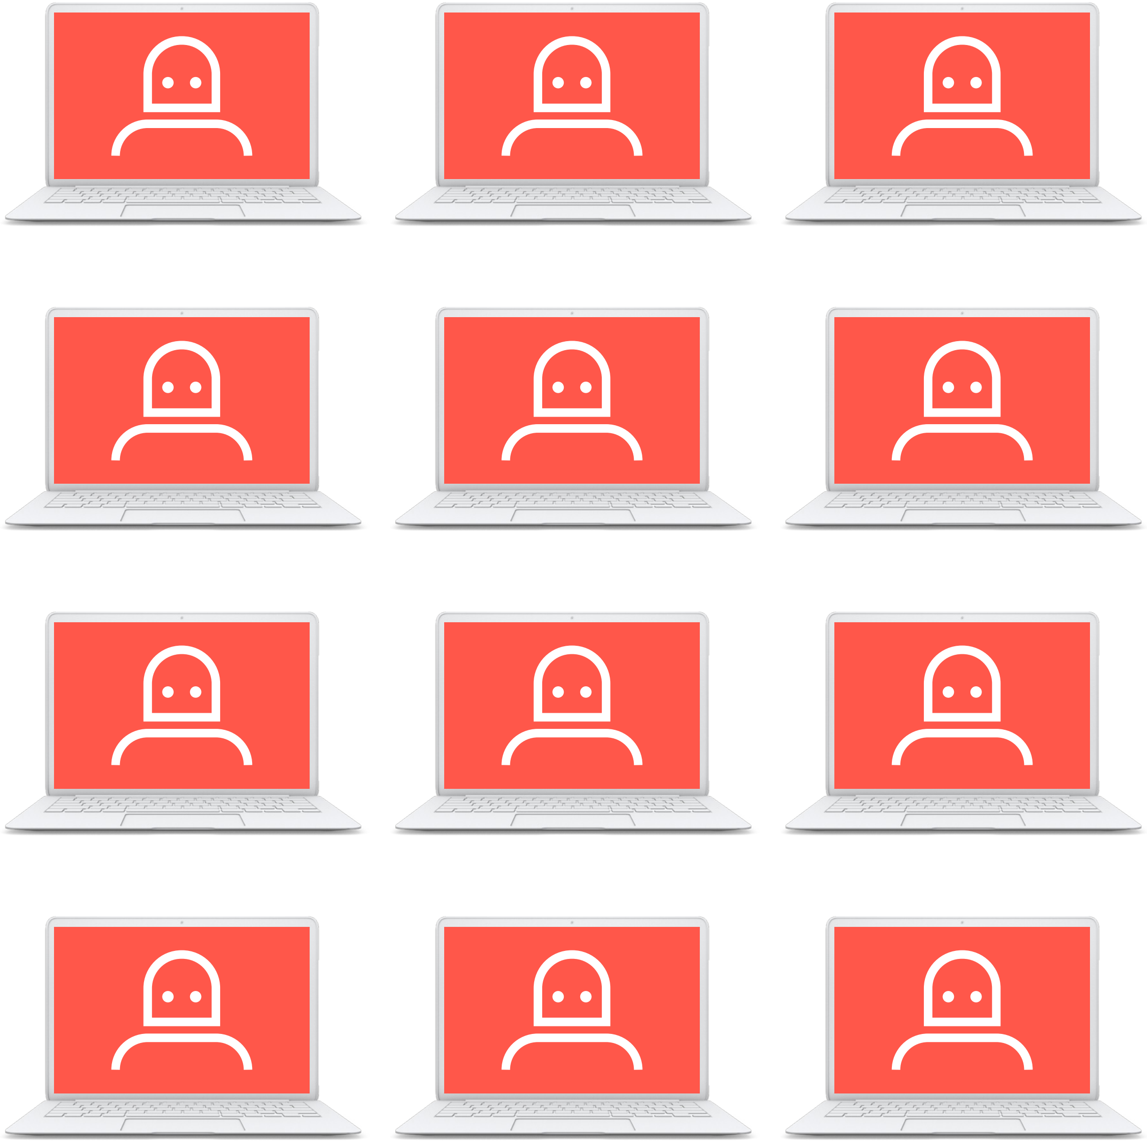

Bots can also be knowingly used as part of a volunteer botnet.

In other words, an everyday computer user will knowingly allow their computer to automatically browse content, in order to get reciprocal traffic to their own content. For example, if someone wanted a higher number of views for their blog, they would join the volunteer botnet to drive up traffic to their site and the sites of other people within the volunteer network.

It's an "I scratch your back; you scratch mine" scenario.

For increased scale or greater control, fraudsters may rent computers from datacenters.

These contain thousands of computers that are available on an hourly, daily, or longer basis; fraudsters rent these computers to commit their crime. And unlike hijacking consumer PCs, using datacenters is completely legal.

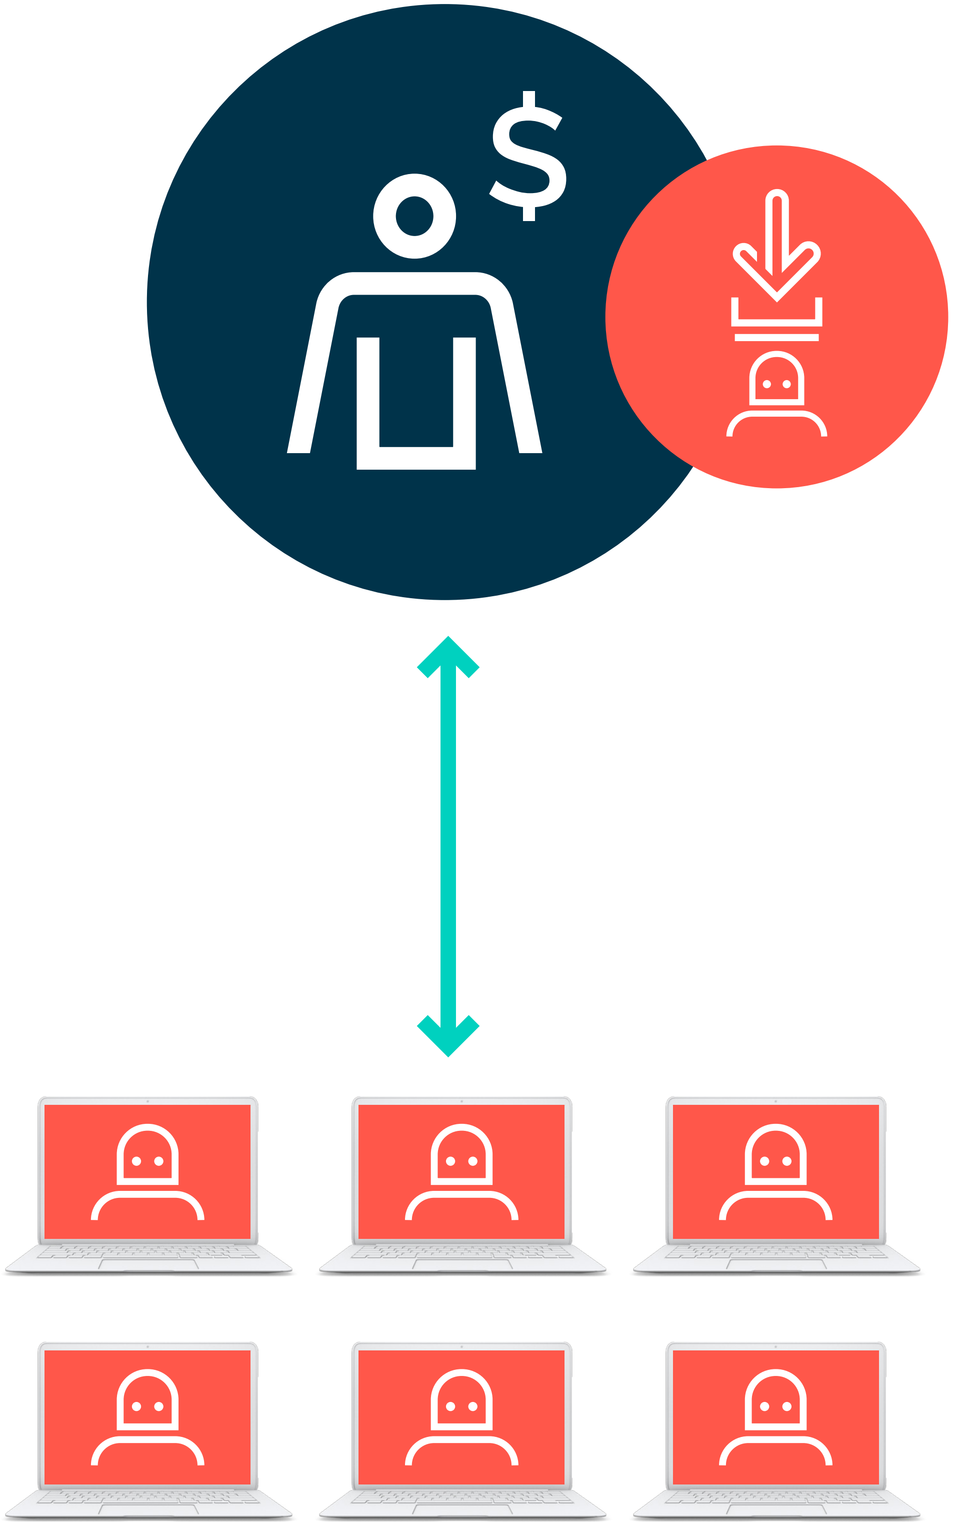

Types of bots

While only 43% of the industry said they understand how fraud is detected, there is increasing demand for transparency when it comes to fraud reporting, especially for bot traffic. Advertisers and media partners need more informative conversations about fraud in order to mitigate the risk within campaigns, which requires more information in general.

There are many bots out there, becoming more sophisticated, and acting more and more like humans. Here's a quick look at two specific bots—Poweliks and Avireen—to give you a general sense of how different bots can operate.

Poweliks

Poweliks is a botnet used for committing both impression fraud and click fraud by generating traffic to ad-supported websites and faking human interaction. It's Internet Explorer-based malware, with a sophisticated human-interaction module and the capability to execute multiple hidden browsers on a single computer (up to 25 at a time!) to simultaneously conduct fraud. It simulates content-specific interactions, can detect videos, hyperlinks, and search bars, and executes applicable actions. It also simulates mouse movements, hovers, and clicks. It has been specifically programmed to counter particular detection mechanisms implemented by ad tech companies.

To learn more about Poweliks, download the white paper written by David Wells, our senior malware analyst.

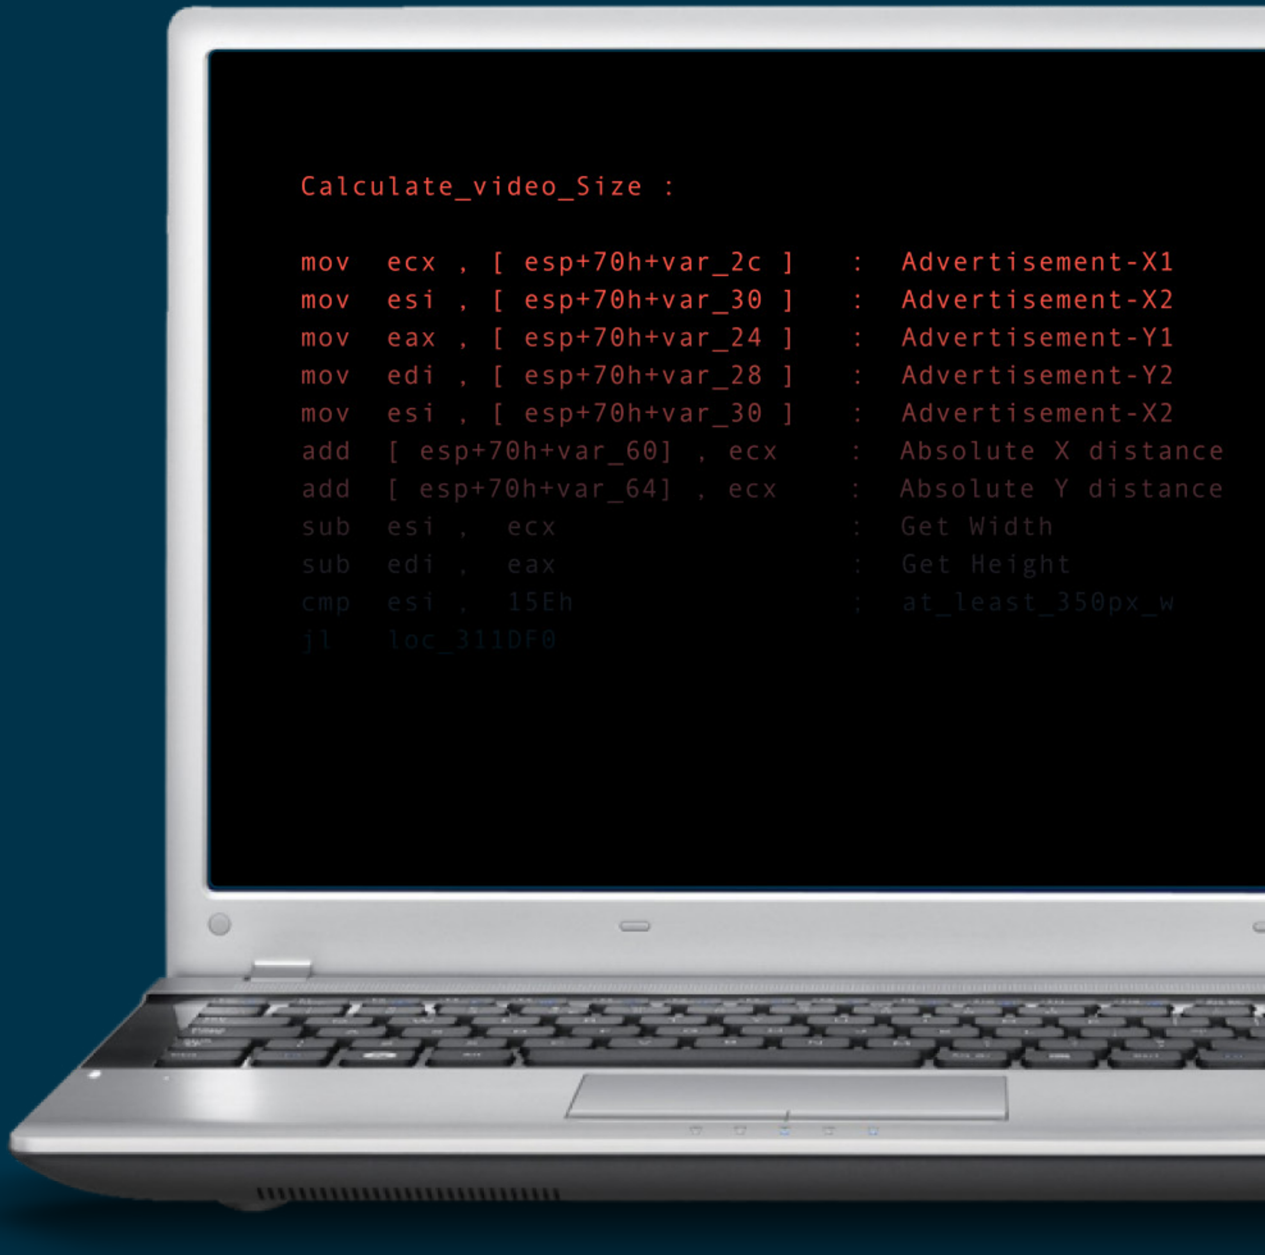

```
Calculate_video_Size :

mov  ecx , [ esp+70h+var_2c ]   :  Advertisement-X1
mov  esi , [ esp+70h+var_30 ]   :  Advertisement-X2
mov  eax , [ esp+70h+var_24 ]   :  Advertisement-Y1
mov  edi , [ esp+70h+var_28 ]   :  Advertisement-Y2
mov  esi , [ esp+70h+var_30 ]   :  Advertisement-X2
add  [ esp+70h+var_60] , ecx    :  Absolute X distance
add  [ esp+70h+var_64] , ecx    :  Absolute Y distance
sub  esi ,  ecx                 :  Get Width
sub  edi ,  eax                 :  Get Height
cmp  esi ,  15Eh                :  at_least_350px_w
jl   loc_311DF0
```

Avireen

Avireen is a bot used for conducting both impression fraud and click fraud, part of a larger malware system called Andromeda, which also engages in other types of malicious activity, including ransomware. It is able to control both Chrome and Internet Explorer, depending on what the user has installed. It impersonates human behavior, simulating mouse movements and hovers. In addition, it leverages the existing user's cookie cache to masquerade as a believable, real human user while doing its browsing and deletes any new cookies that could flag it as a bot.

To learn more about how we discovered Avireen and how it operates, check out our white paper, written by senior malware analyst David Wells.

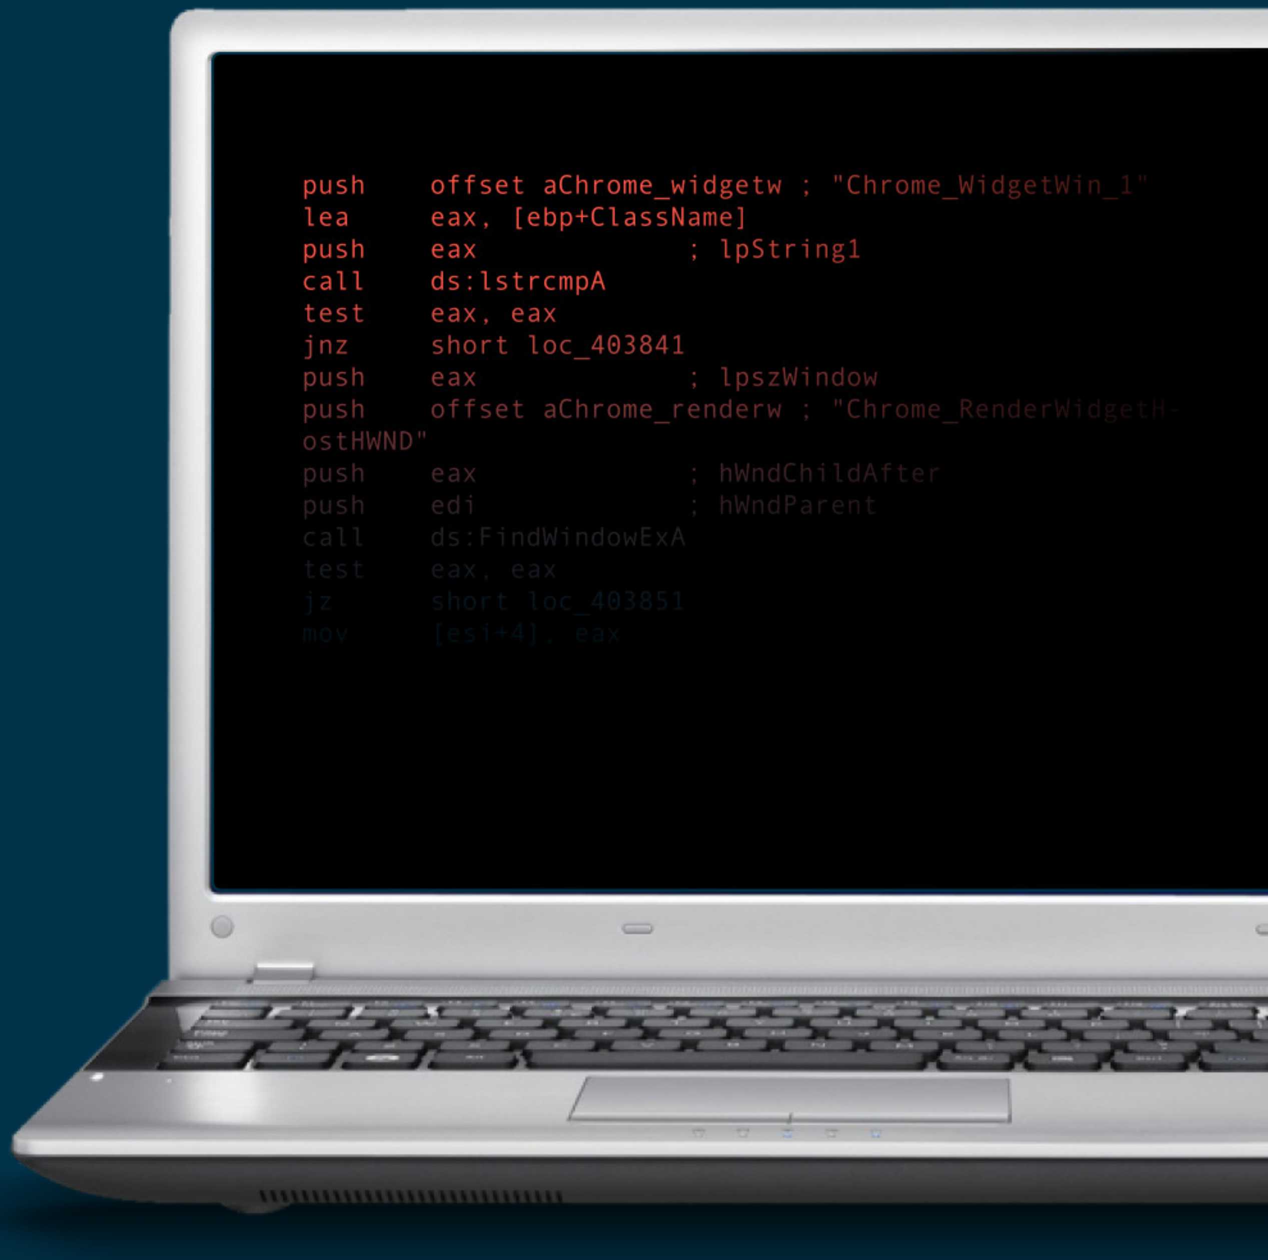

```
push    offset aChrome_widgetw ; "Chrome_WidgetWin_1"
lea     eax, [ebp+ClassName]
push    eax                ; lpString1
call    ds:lstrcmpA
test    eax, eax
jnz     short loc_403841
push    eax                ; lpszWindow
push    offset aChrome_renderw ; "Chrome_RenderWidgetH-
ostHWND"
push    eax                ; hWndChildAfter
push    edi                ; hWndParent
call    ds:FindWindowExA
test    eax, eax
jz      short loc_403851
mov     [esi+4], eax
```

Key takeaways

1

Bots prey on higher-value media

2

Bots can consume ads at any stage of the digital-advertising chain

3

The majority of bots come from residential Internet addresses

4

Bots account for a greater proportion of traffic at night

5

Bots often fill hard-to-reach demographic quotas

6

Bots are acting more and more like humans, able to hover, scroll, and more

CHAPTER FOUR:

Fraud and video

Video ad spend is projected to reach $10.3B in 2016 in the U.S. alone, according to eMarketer.

That's a 34.1% increase from 2015.

Nothing is more powerful than connecting with consumers through sight, sound, and motion.

But video inventory is particularly susceptible to fraud—across premium and programmatic video—because the medium has the highest CPM and the greatest expected impact.

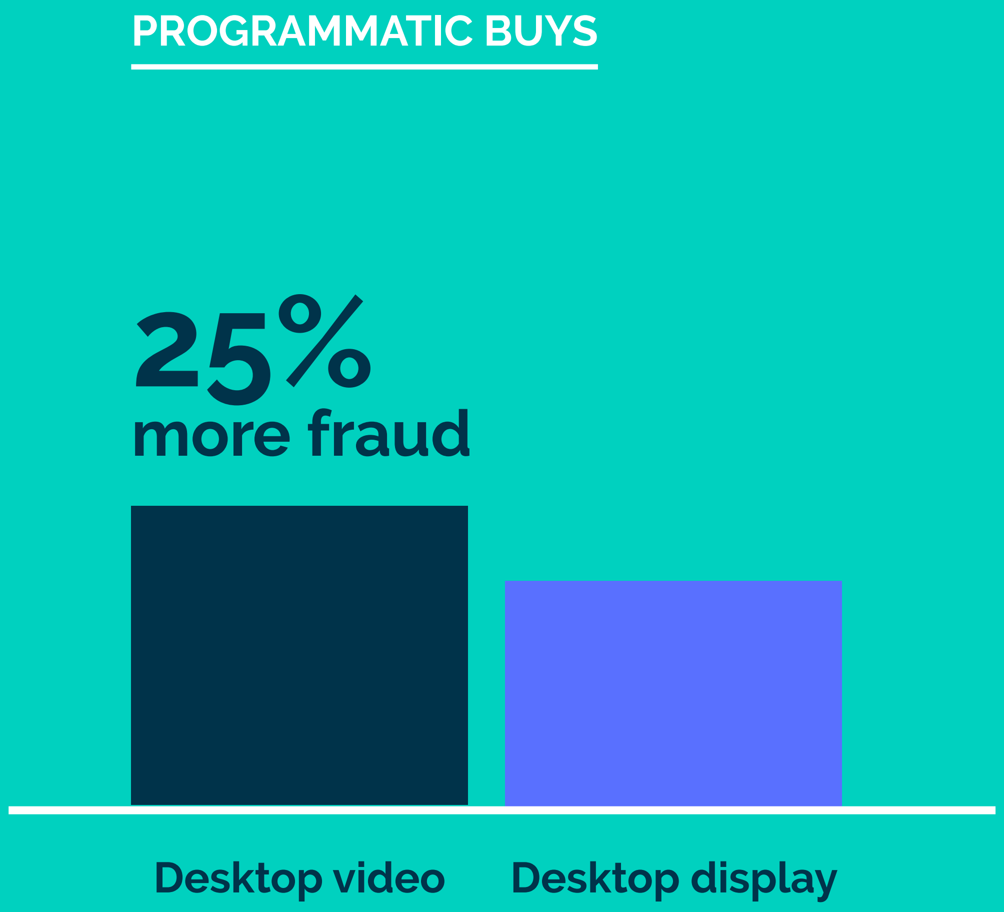

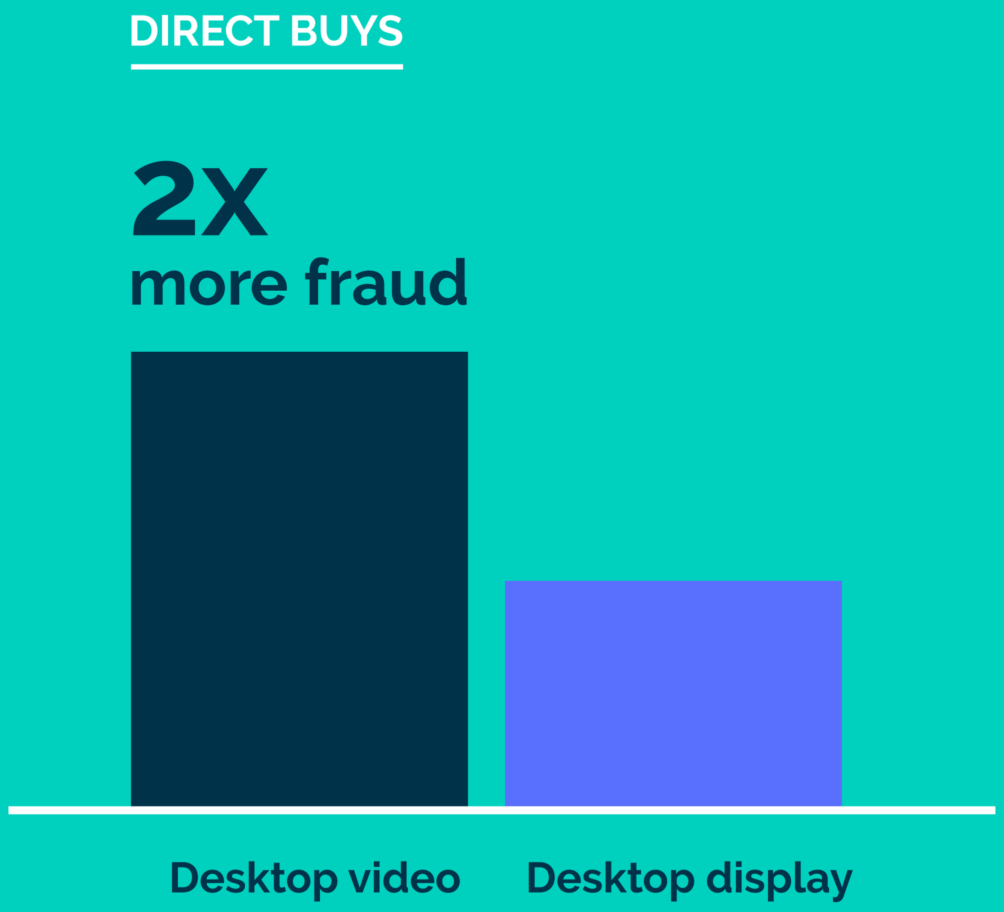

The video-heavy revenue models of many premium publishers make them more susceptible to fraud.

While video fraud often acts the same as display fraud, there are a few key things to note:

Volunteer botnets often focus on individual videos on user-generated content sites where uploaders can share in the revenue.

A site that buys traffic will often start by buying for video, because of that higher ROI.

CHAPTER FIVE:

Fraud and mobile

Mobile ad spend is projected to top $100 billion worldwide in 2016, according to eMarketer.

That's 51% of the entire digital market.

As mobile continues to grow in consumer usage—and as advertising follows—it's expected that fraud techniques will become more tailored, and more pervasive.

Approximately 10% of all display impressions in the U.S. are fraudulent; there is reason to believe that this is equally the case—or more so—in mobile. Increasing evidence suggests that ad fraud is widespread in mobile activity, impacting SSPs, DSPs, brands, agencies, and publishers.

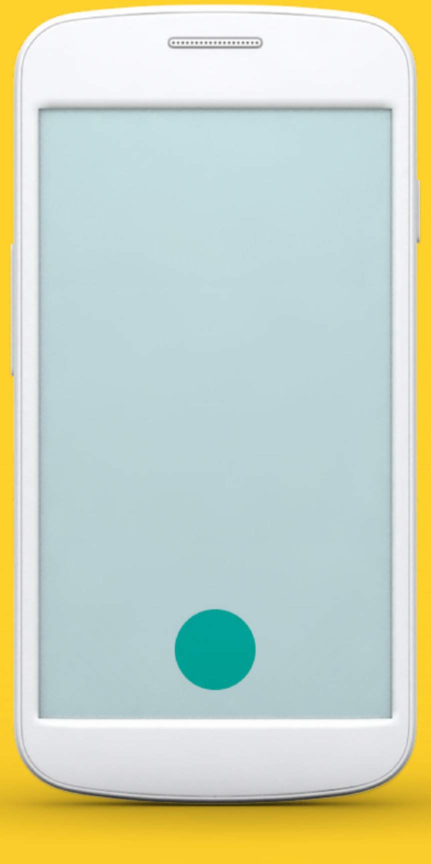

Fraud in mobile can:

Siphon money away from legitimate producers

Distort marketing results

Make effective optimization confusing

Fraud in mobile is not typically caused by bots. While bots can still wreak havoc on mobile, the more prominent type of fraud in mobile 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.

Types of mobile fraud

Mobile fraud can occur in a couple of ways:

1. Within an app
2. In the cloud

Malicious apps

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

1

Background services

Services running in the background are able to render ads even when the app is closed—or not even started! 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 by itself on every device reboot or is triggered by a common activity such as screen lock or volume change.
- Spoofs geo-location/device ID/app name in order to obscure its abnormally high ad volume per device, and remain inconspicuous to ad tech.
- It can even generate clicks on the invisible ads through injected JavaScript.

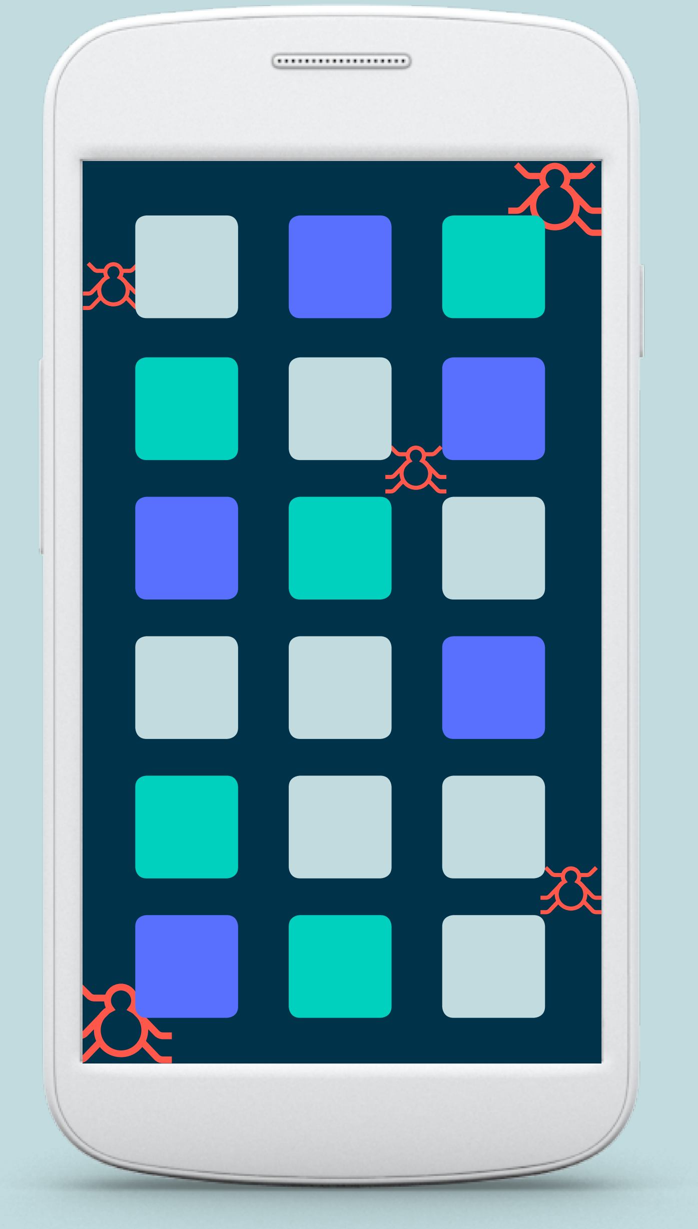

2

App-name spoofing

Similar to domain spoofing in display, apps can submit a false ad-unit identifier or app identifier to the bidding platform. This interferes with:

- Brand safety and contextual targeting
- Detection of apps utilizing background services to load ads

3

Hidden ads

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

- Ad stacking
- Invisible banners

Many apps are used without the user being present, for example, radio apps, GPS, alarm clock apps, etc. These apps must be started by the user, but are designed to remain on for a long duration without the user being present. However, ads can still be shown in these apps.

While this may not be fraudulent, such apps represent extremely low-quality inventory, which should be kept in mind when buying.

74°F Sunny, 3 MPH
9:21
AM
05
WED
DISPLAY AD

Cloud hosting

Mobile devices hosted on Virtual Private Servers (VPS)

In-app impressions are generated by hosting mobile-operating-system emulators or devices in the cloud and running apps that display ads, thus generating ad revenue for the app creators. Having full control over changing the device ID/Geo-location/User-agent can mimic a variety of different app users to ad networks.

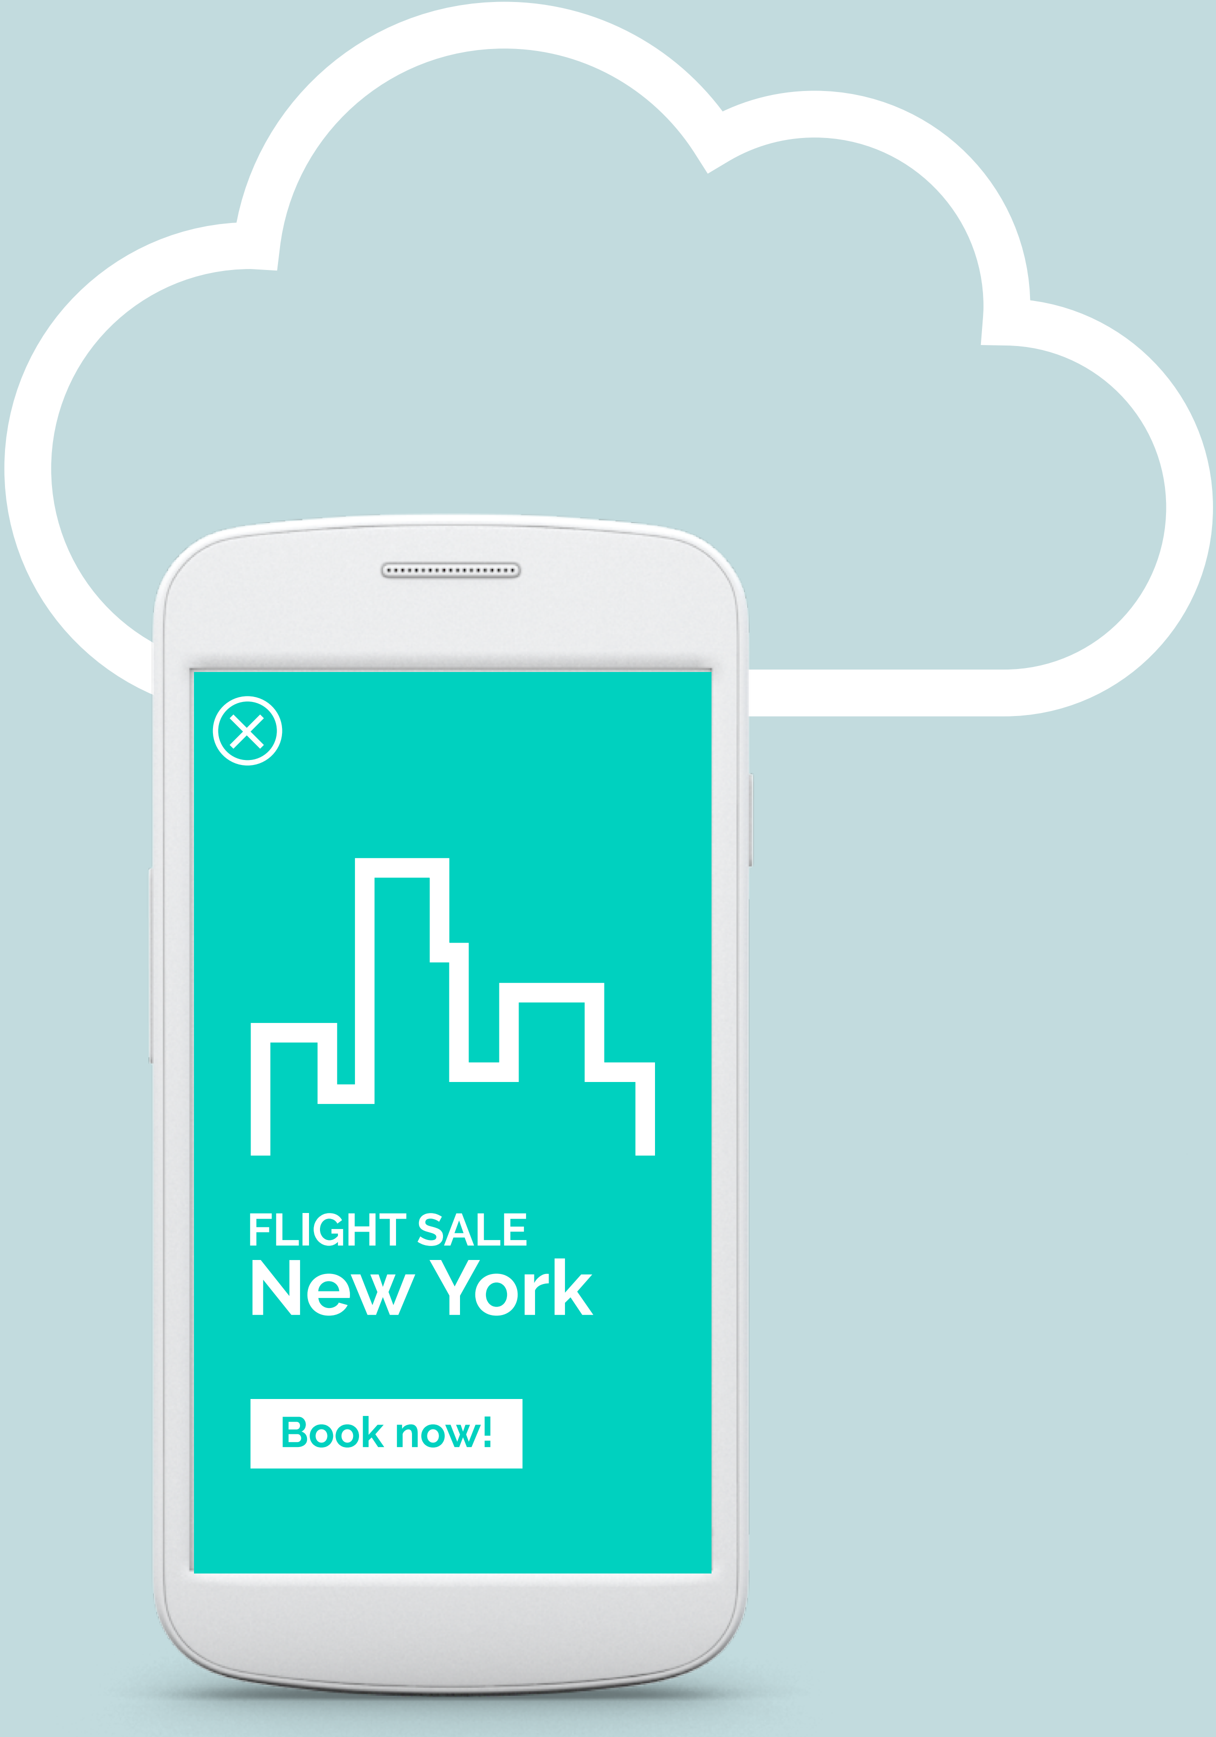

What makes fraud in mobile particularly challenging?

Most types of fraud in mobile advertising are similar to those 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, improving mobile viewability standards and mobile fraud measurement will help reduce its impact.

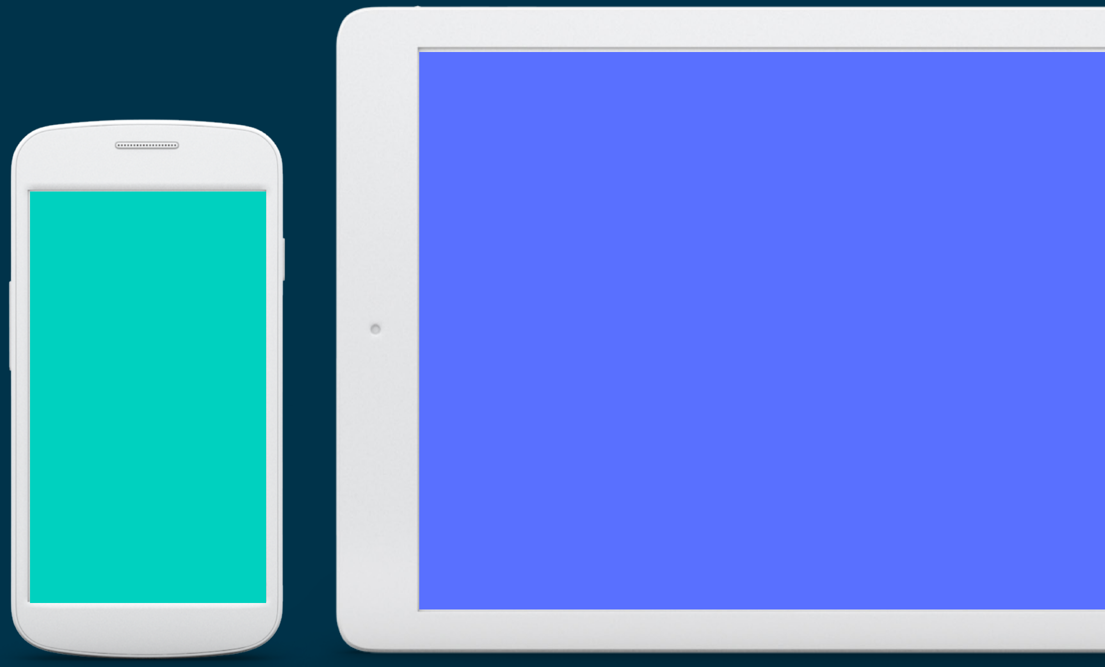

CHAPTER SIX:

Fraud and programmatic

It's no surprise that programmatic digital display ad spending is projected to increase by 44.2% in 2016, and programmatic digital video ad spending by 106.3%, according to eMarketer. After all, programmatic offers a lot of great opportunities and efficiencies.

But it's more likely to contain fraud—up to 4x higher. That's just a fact. The nature of programmatic makes it easier to conceal or lie about the people involved, the quality of the inventory, etc.

What can you do?

☑ If it’s too good to be true, it probably is. Don’t focus as much on low CPMs and CPCs, but instead focus on real KPIs based on your goals (sales, sign-ups, etc.).

☑ Vet your vendors and partners – and their vendors and partners.

☑ Try to gain more visibility and transparency into where the programmatic advertising is being served.

CHAPTER SEVEN:

How to identify and fight fraud

As fraud becomes more sophisticated, the digital industry needs more sophisticated fraud detection to evaluate the legitimacy of impressions and to prevent the buying and selling of fraudulent inventory.

Watch Scott Knoll, our C.E.O. and president, discuss the global fight against fraud with The Drum.

How can the industry fight illegal traffic?

Ad fraud is a complicated phenomenon that involves hackers, different software black markets, traffic brokers, and publishers with varying degrees of awareness of what is happening. Not all aspects are explicitly illegal, and those that are typically occur in countries with indifferent or ineffective cybercrime law enforcement. As a result, **proactive measures for avoiding fraud are required**, rather than relying on criminal-justice systems.

Advertisers that choose to use **blacklists**, which prevent ad delivery to sites that have had a history of fraud, apply a reactive method that immediately shuts down a fraudulent supply channel. However, lists are often not updated frequently enough, and can significantly impact scale. Additionally, premium publishers can fall victim to fraud, even when following best practices, and being placed on a blacklist incorrectly penalizes them. Meanwhile, as soon as a site is blacklisted, a new one can be registered and used to continue the fraud.

There are really three pillars in dealing with ad fraud:

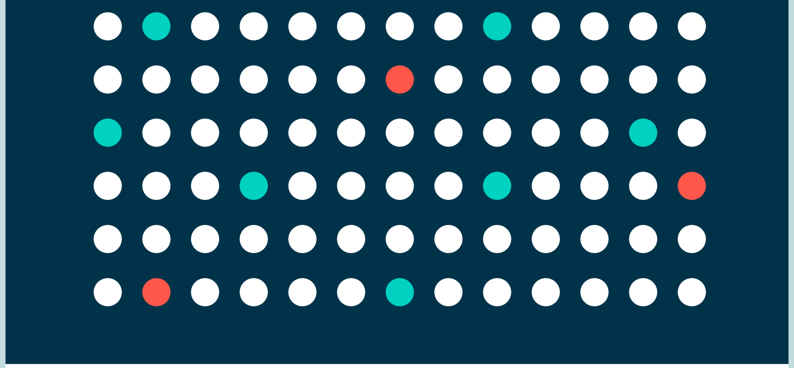

Behavioral and network analysis

Using data science to understand users

Browser and device analysis

Using web technologies to understand implementation

Targeted reconnaissance

Using malware analysis, software disassembly, and the infiltration of hacker communities (also known as black-hat monitoring) to guide detection development and identify emerging threats

These techniques are all required for a well-rounded, sophisticated, program of detection and prevention.

In order to effectively combat fraud, it's critical to develop techniques leveraging data science and advanced web development, both guided by intensive intelligence gathering. Techniques relying on specially designed data collection within the ad display environment are sometimes referred to as session-based signals or side-channel analysis.

With these three pillars as a foundation, today's technology applies advanced learning about fraud to real-time signals to make a decision about the existence of fraud on a given web page.

Behavioral and network analysis

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. Each individual impression can be measured against:

- Delivery channel, web page, inventory source, and user
- Temporal and historical browsing patterns – trends over time
- Geographical distribution – where traffic comes from
- Page interaction – scrolling, clicking, mouse movements

PROS	CONS
• Ensures broad coverage of disparate threat types • Enables fast reaction to emerging threats • Resilient to attempts at circumvention because the methodology is hidden server-side	• Decisions may not be available at impression time, depending on implementation • Requires a massive footprint across billions of impressions in order to yield stable, actionable intelligence • Requires significant processing abilities • Requires expertise in statistics and/or machine learning

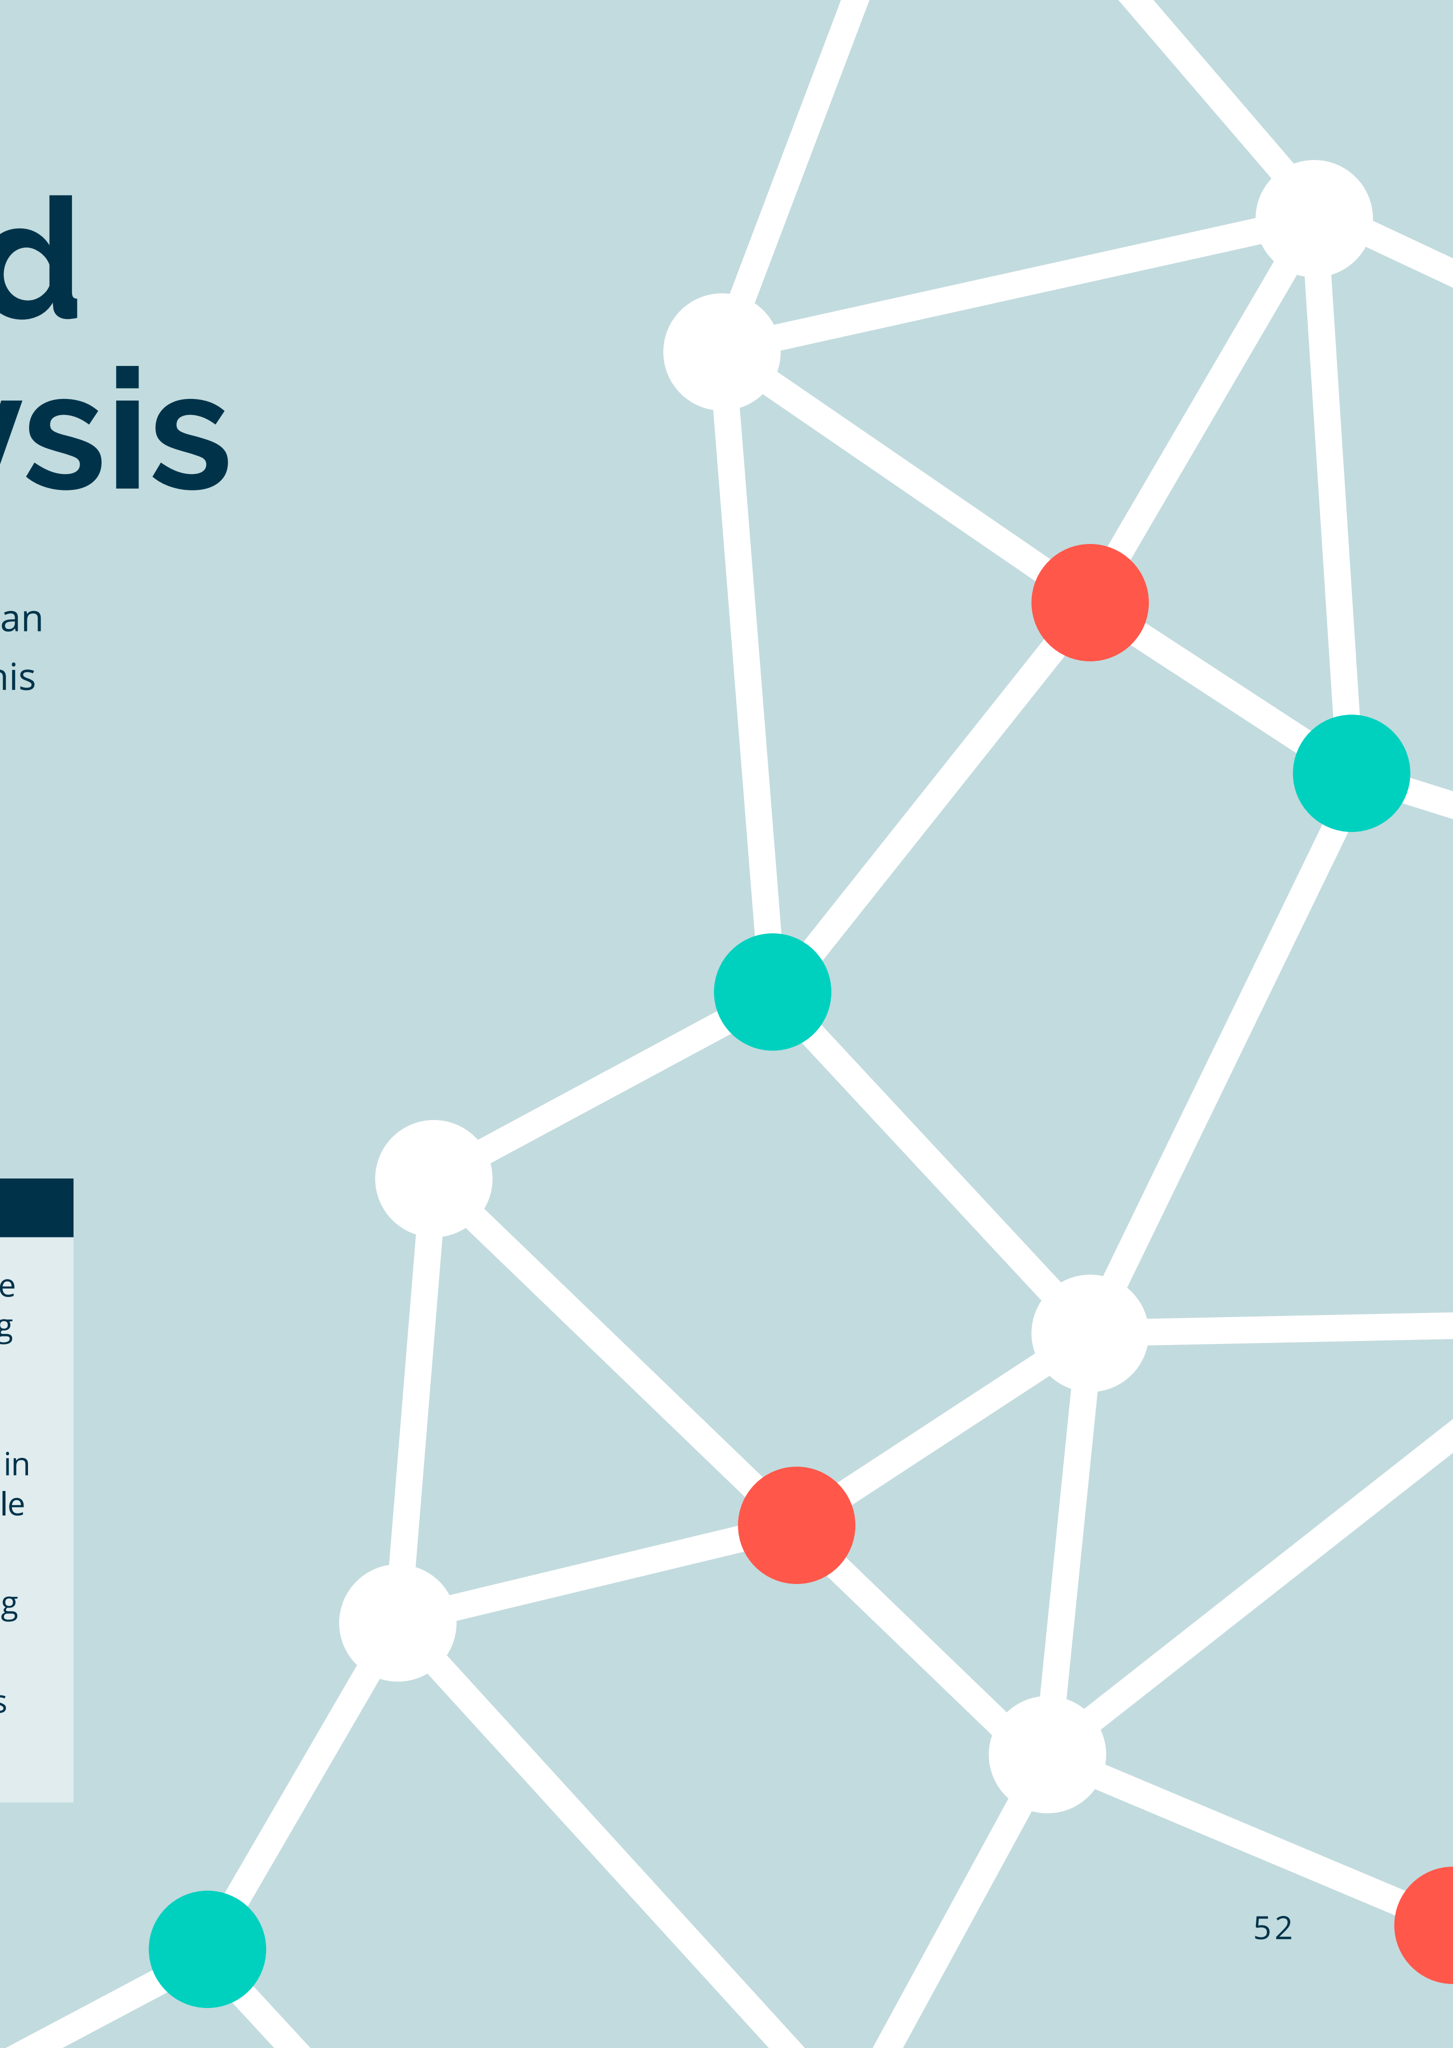

Browser and device analysis

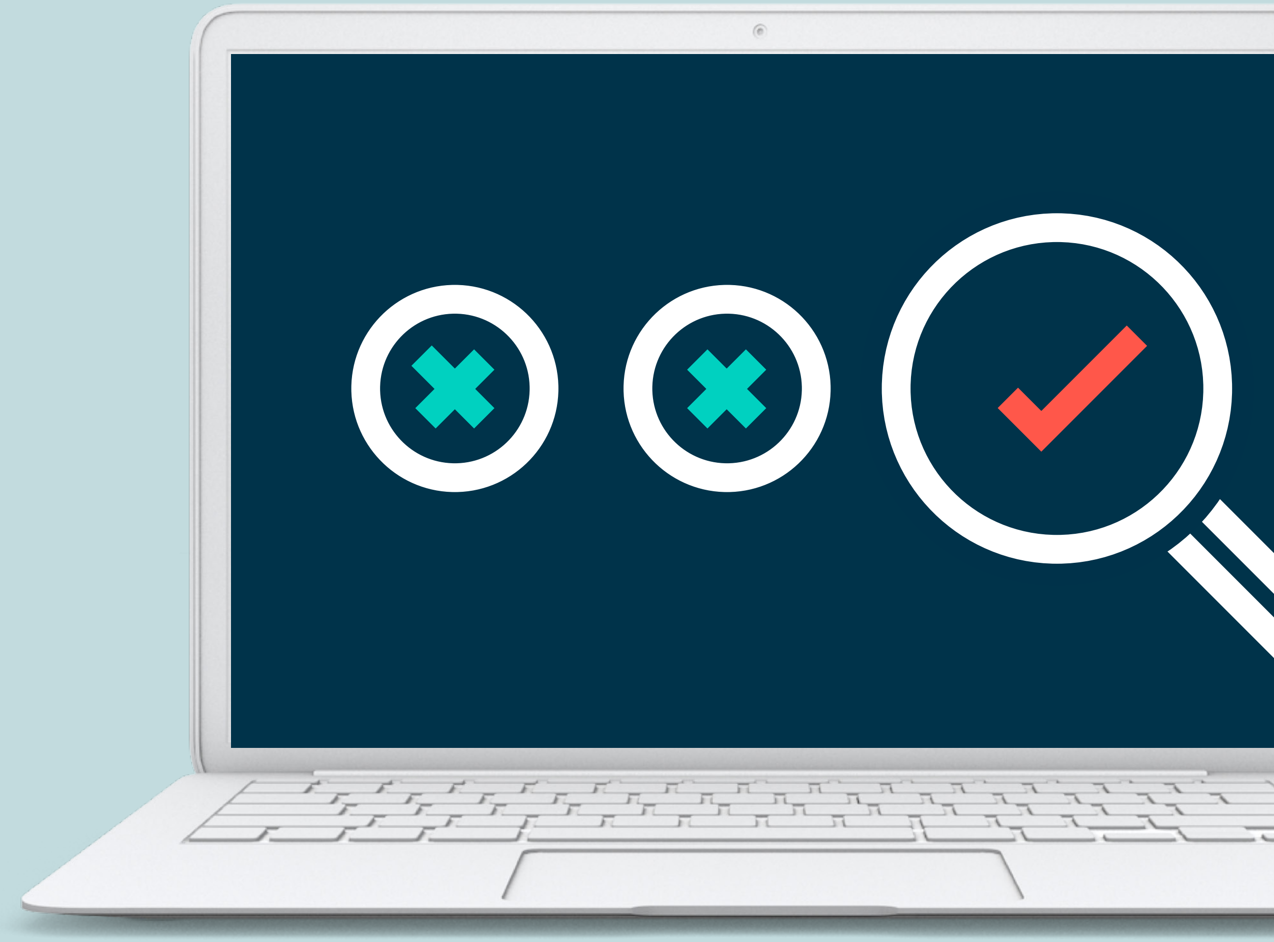

Each bot, whether a sophisticated strain of malware or a DIY script, has a signature set of characteristics that can be identified through detailed mapping of the browser environment and device characteristics. A machine infected with a bot will perform differently than an uncompromised computer, despite the fact that the abnormality may not be apparent to the user. Additionally, inspection of page layout and the ad-serving chain can reveal hidden ads, domain spoofing, and more. Signals useful for this type of analysis include:

- Browser support of common features
- Hardware utilization
- Discrepancies between viewability measurement techniques
- “Honeypots” to induce a bot to interact with the page in a way a human could not
- Distinctive page styling
- Bot-like utilization of the mouse, scrollbar, etc.

PROS	CONS
• Highly diagnostic • Guarantees complete coverage for an individual bot • Instant action, with no wait for statistical significance	• Could be unstable if the bots are modified • Methodology is exposed client-side • Requires expertise in malware and browser technologies • Limited to browser environments that use JavaScript

Comparing the two methods

Fraudsters are actively attempting to evade detection, making the situation ever-changing and more complicated. The most robust solutions are those that can combine these two methodologies.

Behavioral and network analysis		Browser and device analysis
Greater coverage	Accuracy	More precise
Quick response	New threats	Complete removal
Secure methodology	Risk/rewards	Instantaneous

Malware analysis and targeted reconnaissance

Fighting fraudsters is an arms race, but the good news is the advertising industry has acknowledged the problem and taken on the challenge. In particular, ad tech companies are increasingly engaging with the cybersecurity community to determine the most effective ways to infiltrate hacker communities and discover threats to advertisers. Just as fraudsters can try to reverse engineer security signals from tech companies, malware analysts can reverse engineer bots and other forms of fraud through activities such as:

- Disassembly of fraudulent malware and software
- Direct analysis of paid traffic
- Infiltration of hacker communities
- Social engineering tactics

Continued, dedicated, research and development of anti-fraud technology is the final critical piece to solving the ad fraud problem.

CHAPTER EIGHT:

How to protect yourself from 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.

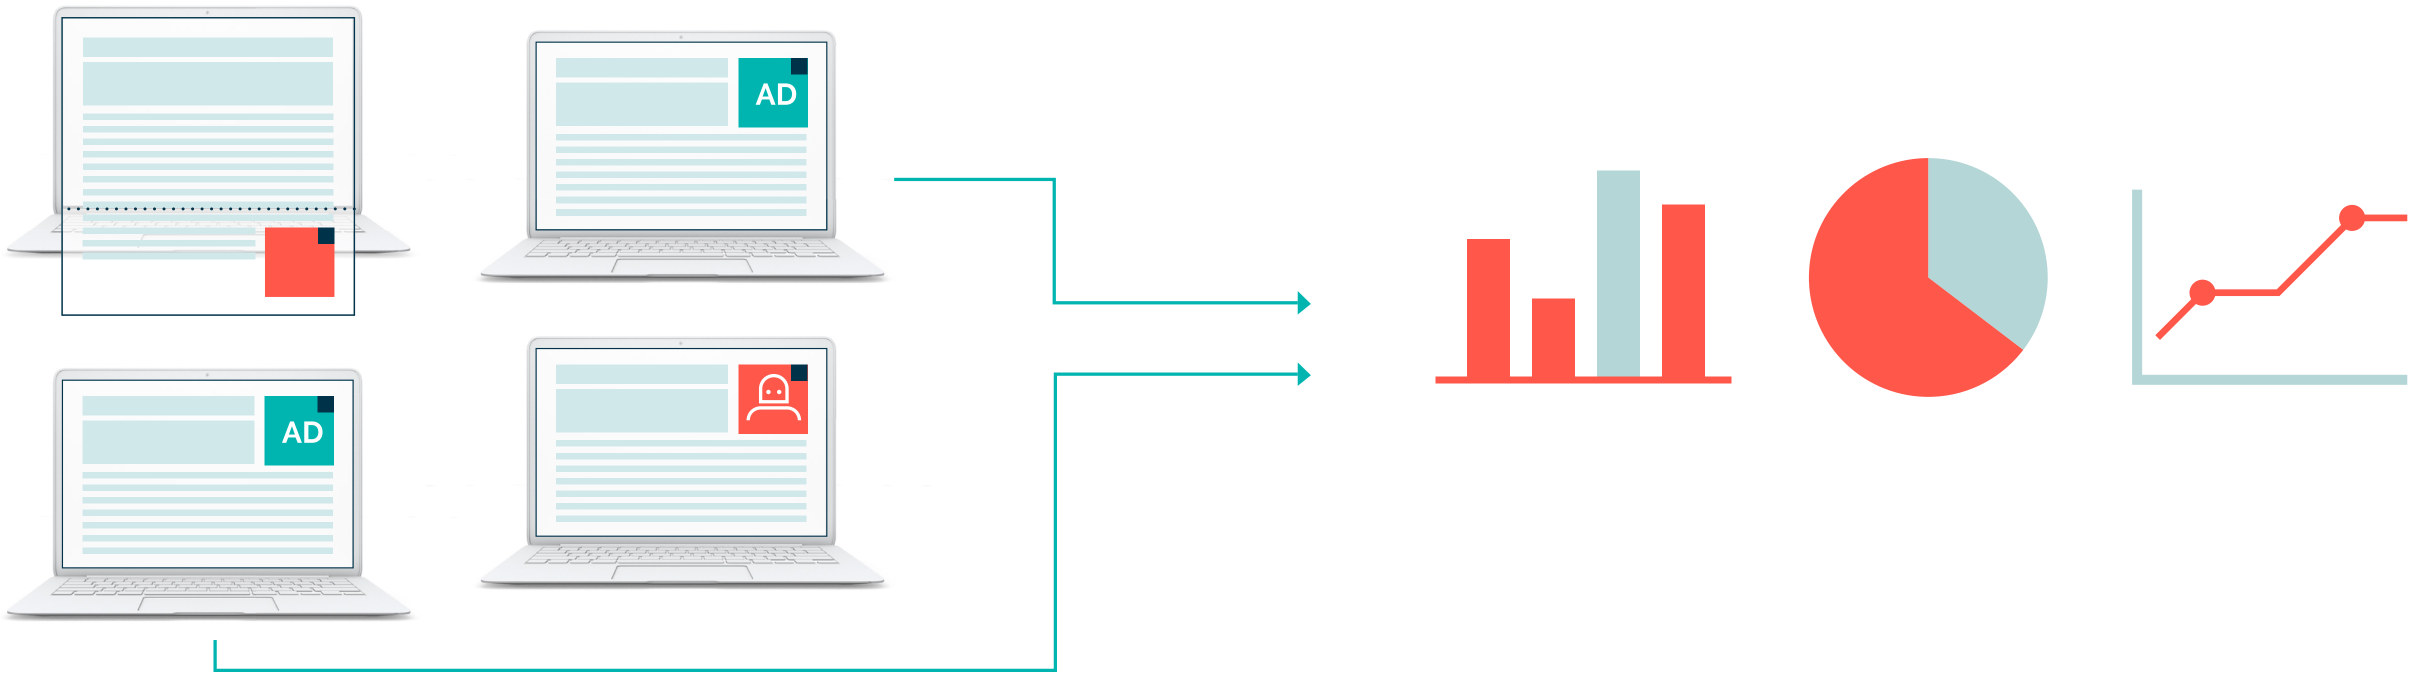

Identify fraud

☑ Measure fraud across all campaigns to understand aggregate performance against fraud.

☑ Use fraud solutions that have been accredited by the MRC for both general and sophisticated IVT.

☑ Follow the MRC guidelines for IVT detection and filtration.

☑ Ask your ad server, fraud solution, or other vendor how it measures for bots and other forms of IVT.

☑ Offer and request more transparency into inventory and traffic, including sourced traffic and audience extension

☑ Use verification and fraud services that can confirm ads were delivered on plan (to the sites, devices, geographies, and audiences desired); whether the environment had ad clutter and other placement concerns; whether it was brand-safe.

Prevent fraud

- ☑ **Block fraudulent impressions before they hit the creative ad server.**
- ☑ **Anti-target infected machines that have been tagged in order to prevent future ad targeting.**
- ☑ **Anti-target pages that have historical levels of fraud, which can be tracked through page-level scoring.**
- ☑ **Use blacklisting and/or whitelisting.**
- ☑ **Use pre-bid screening.**

The biggest factor here is awareness and participation. It's critical that all members of the digital ecosystem are a part of the process and solution.

High level steps

- ☑ Budget for security and fraud solutions.
- ☑ Vet your vendors and partners, and their vendors and partners – and everyone's technology solutions.
- ☑ Consider adding language around this issue into terms and conditions of agreements, RFPs, IOs.
- ☑ Stay on top of industry initiatives; support the Trustworthy Accountability Group.
- ☑ Seek and deliver make-goods for IVT.
- ☑ Monitor compliance.

Questions to ask your vendors

Some of the following questions are a bit specific and technical, but any vendor you're considering working with should be able to quickly and easily answer them. What the answers are – that's what matters.

1. Do you invest in both the technology and human resources required to fight fraud?
2. Is your technology built in-house or outsourced?
3. Do you have programmatic anti-fraud solutions and anti-targeting capabilities?
4. Do you look at fraud data from both a holistic and granular perspective?
5. Botnet and web-browser technologies are constantly changing. What sort of internal testing and verification processes are in place to ensure yours is up-to-date?
6. Are you accredited for general IVT detection by the MRC? What about sophisticated IVT?
7. What techniques do you use to identify network addresses that host bots (beyond blocking IP addresses of known data centers)?
8. What techniques are used to distinguish bot browsers from human browsers at the impression level?
9. What processes are in place to verify that fraud-detection techniques are not flagging human users as bots (false positives)?
10. What techniques are used to verify that traffic from ad campaigns meets targeting requirements?
11. Independent of viewability, what techniques are used to verify that ad impressions were properly rendered on web pages?
12. What techniques are used to detect domain spoofing?
13. Can you block ads from serving in real time if fraud is detected?
14. Do you conduct continued research in the identification and prevention of fraud?
15. Do you capture bots and analyze them directly to identify how they work?

For the sell side

It might be advertisers' money on the line, but the sell-side reputation—and potential revenue—is on the line as well! In order to protect yourselves, and keep advertisers spending with you, it's important to take these actions.

☑ **Be vigilant.**
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.

☑ **Transparency, transparency, transparency.**
They're all asking for it; make sure you can supply it. By helping advertisers to monitor their media investments more closely, you'll earn their trust and dollars.

☑ **Pay special attention to video.**
It's the most expensive, and therefore deserves a little extra love.

☑ **Enable measurement and monitoring**
for viewability, engagement, and bot detection.

Taking all of these measures will help maintain your quality media, retain business, and allow you to increase the value of your media.

Conclusion

Fraud is everybody's problem.

It hurts advertisers, who dedicate precious resources to getting their message out in the right way. It hurts publishers, from the premium site whose name is sullied by domain spoofing to the mom-and-pop site whose livelihood is threatened by depressed programmatic CPMs. And fraud adds incalculable friction throughout the ecosystem with disputes and make-goods and confusion.

We are ever-advancing as an industry in our ability to deliver the right ad to the right person at the right time – and make sure it's seen! But fraud is more than just unviewable, mis-targeted, or ineffective inventory; it's a breach of trust. And to continue building the future our customers deserve, we have to work together to restore it. That means getting certified for, at minimum, GIVT filtration, and demanding your partners do the same. That means ensuring that you have a partner accredited for SIVT filtration. That means supporting cooperative programs like the Trustworthy Accountability Group. That means investing in better performance metrics so that a bot generating a network request isn't worth anything.

Looking the other way is no longer an option. No one company can solve it for the rest.

Glossary

Ad injection

Inserting ads into an app, web page, etc., without the consent of the publisher or operator of that resource. The ad can be visible or hidden.

Ad stacking

Placing multiple ads on top of each other in a single placement, with only the top ad being viewable. This is a form of impression fraud because the advertiser is paying for impressions even if the end user is not seeing an ad.

Adware traffic/ad injection

A device where a user is present and additional HTML or ad requests are made by the adware independently of the content being requested by the user. Adware may also contain a function to inject an ad from the software onto a web page as the user browses, rather than the ad being delivered by the publisher of the web page.

Auto-refresh

A page ad unit enabled to request a new rendered asset more than once and at periodic intervals.

Blacklisting

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

Bot

Short for robot; refers to a software program that carries out automated tasks on the Internet. There are good bots and bad bots. They may intentionally or unintentionally view ads, watch videos, click on ads, etc.

Bot detection

The detection and differentiation of bot traffic and bot impressions from human traffic and human impressions.

Bot prevention

The prevention of bot traffic and bot impressions before the inventory is bought or sold.

Bot traffic

Nonhuman traffic designed to mimic users and inflate audience numbers.

Botnet

A group of computers taken over by software.

Browser pre-rendering

A device makes HTML or ad requests prior to expected human-initiated navigation to the requested resources.

Cookie stuffing

A client is provided with cookies from other domains as if the user had visited those.

Datacenter traffic

Traffic originating from servers in datacenters, 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 datacenters while still being delivered to human users.

Domain spoofing

HTML or an ad request that attempts to represent a site, device, etc., other than the actual placement. This tricks advertisers and ad exchanges into thinking the inventory is legitimate. This is also called *domain laundering*.

Hidden ad impressions

Impressions that are not actually seen by people because they are hidden behind other ads or website content (as in ad stacking), 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 HTML or make ad requests that are not under the control of a user and made without the user's consent.

Incentivized browsing

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

Invalid Traffic (IVT)

Also referred to as *Nonhuman Traffic (NHT)* or *Suspicious Activity Detection (SAD)*, it is online traffic generated from machines or other bot activity that interacts with digital ads.

General Invalid Traffic (GIVT)

Traffic that comes from known, nonhuman sources on publicly available IP lists. It can be identified through routine means of filtration. Key examples include datacenter traffic; bots and spiders or other crawlers masquerading as legitimate users; non-browser user-agent headers; hidden, stacked, covered, or otherwise unviewable ads; pre-fetch or browser pre-rendering traffic; and invalid proxy traffic.

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.

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.

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.

Sophisticated bot

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

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. TAG has developed an anti-fraud working group with a mission to improve trust, transparency, and accountability by developing tools, standards, and technologies to eliminate fraud. TAG works to combat the negative impact of fraudulent traffic by:

- Planning to create, maintain, and share the threat list, a database of domains that have been identified as known sources of fraudulent bot traffic for digital ads.
- Developing and enhancing anti-fraud standards and protocols for all types of entities
- Developing tools both to identify fraudulent activity, and to better identify reputable companies in the chain that are not associated with fraudulent conduct.

In May 2015 TAG unveiled its Fraud Threat List, a shared database of domains that are known sources of nonhuman traffic. Shortly thereafter TAG launched the Data Center IP list, which identifies sources of nonhuman traffic based upon IP addresses. Support of TAG's initiatives is a crucial step in creating a transparent and legitimate digital advertising ecosystem. Every company across the ecosystem should register with TAG in order to ensure they are doing business with trusted partners.

About IAS

Integral Ad Science (IAS) is a global technology and data company that builds verification, optimization, and analytics solutions to empower the advertising industry to effectively influence consumers everywhere, on every device. We solve the most pressing problems for brands, agencies, publishers, and technology companies by ensuring that every impression has the opportunity to be effective, optimizing towards opportunities to consistently improve results, and analyzing digital's impact on consumer actions. Built on data science and engineering, IAS is headquartered in New York with global operations in ten countries. Our growth and innovation have been recognized in Inc. 500, Crain's Fast 50, Forbes America's Most Promising Companies, and Business Insider's Hottest Pre-IPO Ad Tech Startups.

integralads.com | @integralads