



Joint Industry Committee for Web Standards

JICWEBS

Traffic Fraud: Best Practices for Reducing Risk to Exposure

Version 1 Issued June 2015

CONTENTS

1.	Introduction	2
2.	Executive Summary	2
3.	Overview	3
	3.1 How traffic bots infect legitimate systems	3
	3.2 How traffic bots generate false traffic	3
	3.3 How traffic fraudsters make money	3
	3.4 Why you should care	4
	3.5 What you can do	4
4.	Best Practices for Buyers	5
	4.1 Set goals	5
	4.2 Manage the relationship	5
	4.3 Measure results	5
	4.4 Address traffic fraud	6
	4.5 Questions to ask publishers	6
5.	Best Practices for Publishers	7
	5.1 Questions to ask inventory sources	7
6.	Best Practices for Networks	9
	6.1 Take notice	9
	6.2 Make a stand	9
	6.3 Address the bad actors	10
7.	Closing	10
8.	Appendix	11

1. Introduction

This document was originally created by the IAB US Traffic of Good Intent Taskforce, and has been revised and adapted by the UK Anti-Fraud Commercial and UK Anti-Fraud Technical Working Groups¹.

The JICWEBS Cross-Industry Anti-Fraud Working Groups were established in December 2014 with the purpose of creating cross-industry guidelines and principles to educate the wider market to reduce the risk of exposure to ad fraud, creating a safer, more transparent supply chain

2 Executive Summary

Advertisers expect that all online content is delivered to human audiences. However a portion of digital advertising is being diverted by nefarious entities that exploit the ecosystem to deliver fraudulent traffic.

The potential for fraud exists anywhere that media spending is significant and performance metrics are ambiguous or easily gamed. Unfortunately, online advertising is vulnerable to this. Nefarious groups have found ways to profit from infiltrating legitimate systems and generating false ad views, ad clicks and site visits using robotic programs.

Robotic traffic - known popularly as “bots” - is driven by code and not humans. These bots are often smart enough to mimic human behavior, and they can be difficult to detect. While more sophisticated bots can simulate conversions such as clicking through to sites, they don't generate real conversions by buying goods and services, and they certainly don't engage with brands. Activity generated by these bots waters down engagement metrics driven by human traffic, which dilutes the value of legitimate publisher inventory. Advertisers end up spending money on specious ad impressions never seen by humans. In addition to bot traffic, other fraudulent activities exist such as URL cloaking where advertising purported to be placed on legitimate sites but are in fact placed on a different site. A taxonomy for common fraudulent traffic is detailed in the UK Fraud Taxonomy for Digital Display Advertising document

Traffic fraud takes an organized effort to generate “results” that fool both buyers and sellers. The fraudsters are not just gaming the system; they are often engaging in other organized criminal activity. Ignoring traffic fraud enriches those engaged in such activities. Identifying traffic fraud and destroying the economics that drive this nefarious activity is imperative to achieving trust in the industry.

This document outlines steps that can be taken by individual businesses in the digital advertising marketplace to address traffic fraud within their organizations. Both the buy-side and sell-side need to play a role in defending against traffic fraud and improving the digital ecosystem.

¹ See Appendix

3. Overview

Traffic fraud is not always easy to detect. Knowing how the bad actors operate can help reduce your risk of being victimized by them.

3.1 How traffic bots infect legitimate systems

Fraudsters often operate undetected in legitimate systems by masquerading as ordinary content and tools typically implemented by unwitting consumers.

Some of the ways they infiltrate systems:

- Getting consumers to install toolbars in their browsers.
- Simulating applications such as games or video players in an alternate browser.
- Bundling hidden applications with consumer downloads.
- Inserting code snippets, undetected, on publisher sites.
- Using malicious code to exploit operating systems and browser security vulnerabilities to surreptitiously install fraudulent traffic-generating code.

3.2 How traffic bots generate false traffic

After infiltrating legitimate systems, fraudsters can use bot code in different ways to generate false traffic. They often operate just under the surface or when human users aren't present to detect foul play.

Some of the ways bot code generates false traffic:

- Generating ad views while consumers browse unaware.
- Hijacking user controls to generate fake clicks when the computer is dormant.
- Running invisible processes behind the scenes to simulate consumer activity.
- Compromising cookie data to simulate high-value consumers.

3.3 How traffic fraudsters make money

Business models used by those perpetrating fraud vary, but a common theme tends to be their high profit margins. Even though fraudulent traffic accounts for only a small percentage of real human traffic, that small percentage generates disproportionate ad inventory and diversion of digital marketing budgets. Some premium publishers are themselves buyers of traffic. If they buy from traffic vendors that are bot-riddled, the bots end up on the premium sites, inflating their impression volume. Some of the ways those undertaking fraud make money:

- Selling cheap traffic to publishers wishing to extend their inventory.
- Selling their own robotic inventory to buyers through an exchange that is unaware

- of the traffic source.
- Becoming part of a legitimate network that pools inventory for buyers. The legitimate network is often unaware of any foul play.
- Creating a network by infecting legitimate sites with bot code, known as a “botnet,” that generates traffic for which they can bill.
- Making ad calls that serve ads one behind another (stacked) or into 1x1 pixel frames, creating hidden ad inventory that generates false impressions from both human and non-human traffic.

Fraudsters contaminate legitimate businesses. For example, purchasing traffic is a generally acceptable way for publishers to extend audience and increase inventory. When legitimate businesses unknowingly purchase traffic from fraudulent businesses, they pollute their available inventory and undermine their relationship of trust with advertisers.

3.4 Why you should care

Allowing the bad actors in our industry to profit from traffic fraud affects the entire online community. In addition to diluting inventory value and diverting funds from legitimate businesses, traffic fraud undermines the integrity of digital media.

Some of the negative impacts of traffic fraud:

- Brands waste money on ad campaigns that are served to invisible inventory.
- Digital media is degraded, and brands look elsewhere for their marketing solutions.
- Ad performance and website visit data are contaminated, undermining analysis.
- Artificial fraudulent inventory floods the market and decreases the value of legitimate (real human) inventory.
- Criminal activity is enabled.
- The industry could be subjected to government oversight, negative press and potentially business-dampening enforcement.

3.5 What You Can Do

The solutions to traffic fraud are not always intuitive. For example, the outright blocking of fraudulent traffic gives information to the fraudsters that helps them blend in better and become more difficult to identify. In addition to the following general guidelines, steps specific to buyers, publishers and networks are outlined in subsequent sections. The following general guidelines can help any online business get started:

- Educate yourself about traffic fraud and the risks that it poses to your business.
- Adopt policies and strategies to identify fraud and mitigate its impact.
- If you are an advertiser, set clear objectives for your media campaigns that focus on the measurement of real ROI, which is difficult for fraudsters to falsify. Measures such as click-through rate, completion rate, and last-touch attribution are easy to game.
- Practice safe sourcing and trust only business partners who have earned trust.
- Implement technology to detect and prevent fraud.
- Filter traffic through vendors who prioritize fraud detection.

4. Best Practices for Buyers

Buyers in online media have much to lose when it comes to traffic fraud. Taking steps to inspect the quality of your buys can go a long way toward preventing fraud in the digital marketplace. The following recommendations should help you achieve quality media buys by identifying and eliminating traffic fraud.

4.1 Set goals

Setting goals before buying media is generally a good practice. Some specific recommendations for buying digital media:

- List specific objectives for your media campaign. Don't leave objectives broad and open to interpretation. Examine whether your goals accommodate fraud.
- Be willing to pay the real price for the media you want. For example, pre-roll video targeted to a specific real human audience with good attention will cost more than linear video ads placed at random simply to increase views.
- Document your goals clearly and get the seller to sign off on those goals. Agree to pay only for results that align with what's documented.
- Don't optimize for cost alone. Results that seem too good to be true probably are.

4.2 Manage the relationship

Trustworthy sellers shouldn't have any trouble backing up their claims for quality media. Keep in mind the following points to help you manage the relationship with your sellers:

- Filter media sellers before you buy. Even after the campaign is running, ensure that your sellers are following through.
- Despite the best efforts of sellers, all fraud cannot be eliminated. Determine the risk you are willing to accept and use that model to discount your media.

4.3 Measure results

Since bots can't engage the way humans can, consider measuring campaign results using more sophisticated metrics that ensure humans are interacting with your ads.

The following measures indicate human interaction:

- Purchases
- Subscriptions
- Verifiable brand survey results
- Validated panels
- Other verifiable engagements

Measures that are easy for bots to fake:

- Ad views
- Clicks
- Click-through rate
- Video completes
- Cookie attribution

4.4 Address traffic fraud

Your internal operations can only go so far to filter out traffic fraud. Look to vendors who specialize in detecting and reducing the more sophisticated causes of traffic fraud.

- License technology specifically developed to discern traffic sources. Brand safety, viewability and placement quality are all fine measures of inventory quality, but they cannot detect the presence of non-human traffic.

4.5 Questions to ask publishers

As you filter for publishers that offer quality traffic, asking questions about how they manage quality control can help you find a good fit for your campaign. The following questions and preferred responses can help you get started finding quality sellers:

Do you have your audience measured by verifiable 3rd party systems?

The publisher should have their audience measured by independent vendors so that you can measure the traffic generated for your ads against an independent benchmark, making anomalies easier to spot. Note, however, that some audience measurement vendors' techniques can be easily fooled by fraudulent traffic. Review vendor methodologies when shopping for a vendor you can trust.

Do you have a clean record with third-party brand safety reports?

Bad site quality is not necessarily correlated with traffic fraud, and high quality sites are not immune to traffic fraud. So screening sites for brand safety is an extra measure that helps ensure sites are involved in efforts to reduce fraud. Along with first-party data, sites should be screened against third-party reports to remove fraudulent and inappropriate environments. There are many different types of brand safety detection and prevention. Investigating third-party methodologies can help you make an informed decision.

How do you determine which impressions are exposed to real humans?

Advertisers want to engage with viewers who are engaged with content, not with users who may have left a web page open accidentally. Each site should have a policy and technical methodology to determine which traffic is generated by real humans. The site methodology should cover how they determine suspect fraudulent traffic, and then flag, investigate and remove it.

How do you assure that ads are served as reported, and that URLs are visible to the advertiser?

Ad opportunities on publisher sites should correspond with the site URL that is reported by first- and/or third party campaign performance analytics.

How do you determine whether ads are auto-initiated or user- initiated?

In video, auto-initiated ads are more susceptible to fraud. Publishers may increase their fulfillment numbers by using auto-play ads when an advertiser specifies they want only user-initiated ads. Publishers need to match the ad interaction requested with the ad interaction fulfilled, and report anomalies.

Do you provide protection from malware?

Websites should provide a safe environment for advertisers and consumers by actively screening for malware. Each publisher should be able to provide information on established approaches.

Are impressions generated by malware redirecting to a site?

Unusually large volumes of traffic and poorly performing placements should be investigated for malicious virus activity. Sites should monitor traffic patterns in real time to recognize anomalies resulting from malware.

5. Best Practices for Publishers

Buying traffic increases the risk profile for a premium publisher. If you want to minimize your risk, you may want to consider not buying traffic from non-organic sources. However, even without buying traffic from non-organic sources, you may have some non-human traffic on your media properties, sent there without your control. For example, consumers who install browser tools or applications may inadvertently open the gate to robotic traffic. In addition, bots may be programmed to browse legitimate sites to build up their targeting cookie pool while avoiding detection.

Despite your best efforts, you may find the need to extend your audience and increase your inventory. For example, if you've committed to delivering more ad impressions than are currently available on your media properties, your choice is to either under-deliver for your advertiser or supplement your inventory with purchased traffic.

Be aware of the risks of purchasing traffic, but if you must increase inventory, the following guidelines can help mitigate your risk:

- As a premium publisher purchasing traffic, pay the higher price to buy quality.
- Look for a natural affinity between your content and the purchased audience.
- Use technology to detect non-human traffic on all of the traffic you are buying.
- Don't lower your standards when performance slips below your goals.
- Know your consultants, and where they are sourcing traffic.

5.1 Questions to ask inventory sources

When you purchase traffic, you put yourself in the buyers' shoes. The following questions are nearly identical to the questions that buyers should ask of publishers. Use the following questions to filter traffic sources that promise to help you increase inventory:

Do you have your audience measured by verifiable 3rd party systems?

The publisher should have their audience measured by independent vendors so that you can measure the traffic generated for your ads against an independent benchmark,

making anomalies easier to spot. Note, however, that some audience measurement vendors' techniques can be easily fooled by fraudulent traffic. Review vendor methodologies when shopping for a vendor you can trust.

Do you have a clean record with third-party brand safety reports?

In video, auto-initiated video ads are more susceptible to fraud. Publishers may increase their fulfillment numbers by using autoplay ads when an advertiser specifies they only want user-initiated ads. Publishers need to match the ad interaction requested with the ad interaction fulfilled and report anomalies.

How do you determine which impressions are exposed to real humans?

Advertisers want to engage with viewers who are engaged with content, not with users who may have left a web page open accidentally. Each site should have a policy and technical methodology to determine which traffic is generated by real humans. The site methodology should cover how they determine suspect fraudulent traffic, and then flag, investigate and remove it.

How do you assure that ads are served as reported, and that URLs are visible to the advertiser?

Ad opportunities on traffic-sourced sites should correspond with the site URL that is reported by first- and/or third party campaign performance analytics.

How do you determine whether ads are auto-initiated or user- initiated?

Autoplay ads are more susceptible to traffic fraud and may be used despite specific requests for only user-initiated ads. Sourced traffic should match the ad interaction requested with the ad interaction fulfilled, and report any anomalies.

Do you provide protection from malware?

Websites should provide a safe environment for advertisers and consumers by actively screening for malware. Each traffic source should be able to provide information on their approach.

Are impressions generated by malware redirecting to a site?

Unusually large volumes of traffic and poorly performing placements should be investigated for malicious virus activity. Sites should monitor traffic patterns in real time to recognize anomalies resulting from malware.

6. Best Practices for Networks

Networks can engage in some key efforts to combat non-human traffic. Differentiate yourself in the marketplace by embracing the practices listed here.

6.1 Take notice

Botnet operators work hard to function under the radar. Malicious players shy away from the light. Your best defense is to look for the telltale characteristics of bad actors.

General red flags:

1. **Publisher has no prior history of substantial traffic**
When a publisher comes to you with traffic built overnight, you can be confident it didn't come from hard work and legitimate content.
2. **High audience overlap between disparate websites**
A handful of common visitors to sports sites and sites on the joys of growing flowers is certainly plausible. However, when a large proportion of the audience for either of these sites visits the other site, while this could be a coincidence, it is potential indicator of fraudulent traffic
3. **Browser stats that are inconsistent with known industry usage stats**
Any publisher touting a billion unique visitors probably does NOT have the attention of one-seventh of the world's population.
4. **Publisher seeks out representation**
Publishers who have earned their traffic get noticed and don't need to ask for representation.
5. **More than four tags on a page**
The more ad tags on a page, the lower the quality of the page. While low quality may mean more traffic fraud, keep in mind that more sophisticated fraudsters will flock to higher quality pages to remain undetected.

Red flags in RTB environments:

1. **"No Bid" reason flag and other automated indicators**
The no bid reason flag in OpenRTB 2.2 is an optional flag a bidder can use to tell the exchange that they are not bidding on the inventory because they believe it to be non- intentional. Both exchanges and bidders are encouraged to create and implement non- intentional traffic detection algorithms.
2. **Negatively target traffic fraud**
Consider using anti-fraud tools in RTB to exclude any bid requests with a high probability for fraud.

6.2 Make a stand

Bad actors will take the road of least resistance. Building collaborative relationships in the industry builds a wall of resistance that turns bad actors away.

1. **Partnerships for rapid and free information sharing**
Collaborative relationships with other supply chain partners, such as exchanges and DSPs, will enable an open channel of communication about bad traffic. Sharing

information openly and quickly builds a defense against bad actors.

2. Vendors that offer fraud and malware detection

Your internal efforts to fight unintentional traffic are the right start. But the more successful you are, the more you'll need help to address sophisticated botnet operators. Certain companies have made fighting fraud and malware the core of their business. Integration with one of these companies sets up a security gate through which you can run all new traffic before selling it.

6.3 Address the bad actors

Taking steps to identify and defend against bad actors goes a long way to improving the value of your network. However, their persistence and changing tactics are designed to poke holes in your efforts. Some ways to combat the bad actors once identified are:

1. Sales disincentive

Your sales teams have the best intentions, but they can unwittingly bring in bad traffic sources. Build in a sales disincentive when that happens, so that sales can focus on finding the higher-grade human traffic.

2. Block suspicious ads

All known traffic fraud should be blocked as soon as it's detected. Suspicious traffic should be monitored, and blocked as soon as it is confirmed to be from a fraudulent source. Many advertisers insist that fraudulent traffic be blocked and refuse to pay for any known traffic fraud.

3. Block payment for fraud

Blocking known traffic fraud subsequently blocks payment to bad actors. When suspicious traffic isn't immediately blocked, fraudsters may demand payment for the impressions that were served. Resist their demands, as fraudsters most likely will back down if asked to prove that their traffic is valid.

7. Closing

For all the well-intentioned industry members who read this guide and adopt its practices, there are fraudsters out there who are also reading this document to discover ways to work around your efforts to thwart them. We have purposely avoided describing detailed strategies here that would help the perpetrators achieve their nefarious ends.

Whatever practices you adopt, you must remain diligent and continue to seek out the expertise to defeat the bad actors that are intent on hijacking the system and robbing legitimate businesses.

Further Enquiries

For further enquiries please contact your representative trade body or JICWEBS directly at info@jicwebs.org

Appendix

The UK Anti-Fraud Commercial Working Group is committee of industry experts across the industry from the following companies:

ABC	Internet Advertising Bureau (IAB)
Association of Online Publishers (AOP)	PHD Rocket
Adloox	Reckitt Benckiser
Forensiq	Santander
FT	Shell
Google	Telemetry
GroupM	Videology
Incorporated Society of British Advertisers (ISBA)	VivaKi
Institute of Practitioners in Advertising (IPA)	Yahoo
Integral Ad Science	

The UK Anti-Fraud Technical Working Group is committee of technical experts predominantly from fraud detection vendors and from companies with a leading interest in fraud detection:

ABC	Microsoft
Adloox	Moat
AppNexus	PHD Rocket
Association of Online Publishers (AOP)	Pixalate
comSxcore	Reckitt Benckiser
Crimtan	Rocket Fuel
Forensiq	Santander
Forensiq	Shell
FT	Sizmek
Google	Telemetry
GroupM	Videology
Incorporated Society of British Advertisers (ISBA)	VivaKi
Institute of Practitioners in Advertising (IPA)	WhiteOps
Integral Ad Science	Yahoo
Internet Advertising Bureau (IAB)	YuMe
Meetrics	

This page is intentionally left blank

Joint Industry Committee for Web Standards

Contact us via info@jicwebs.org

© JICWEBS 2015 www.jicwebs.org