The Bot Baseline: Fraud in Digital Advertising

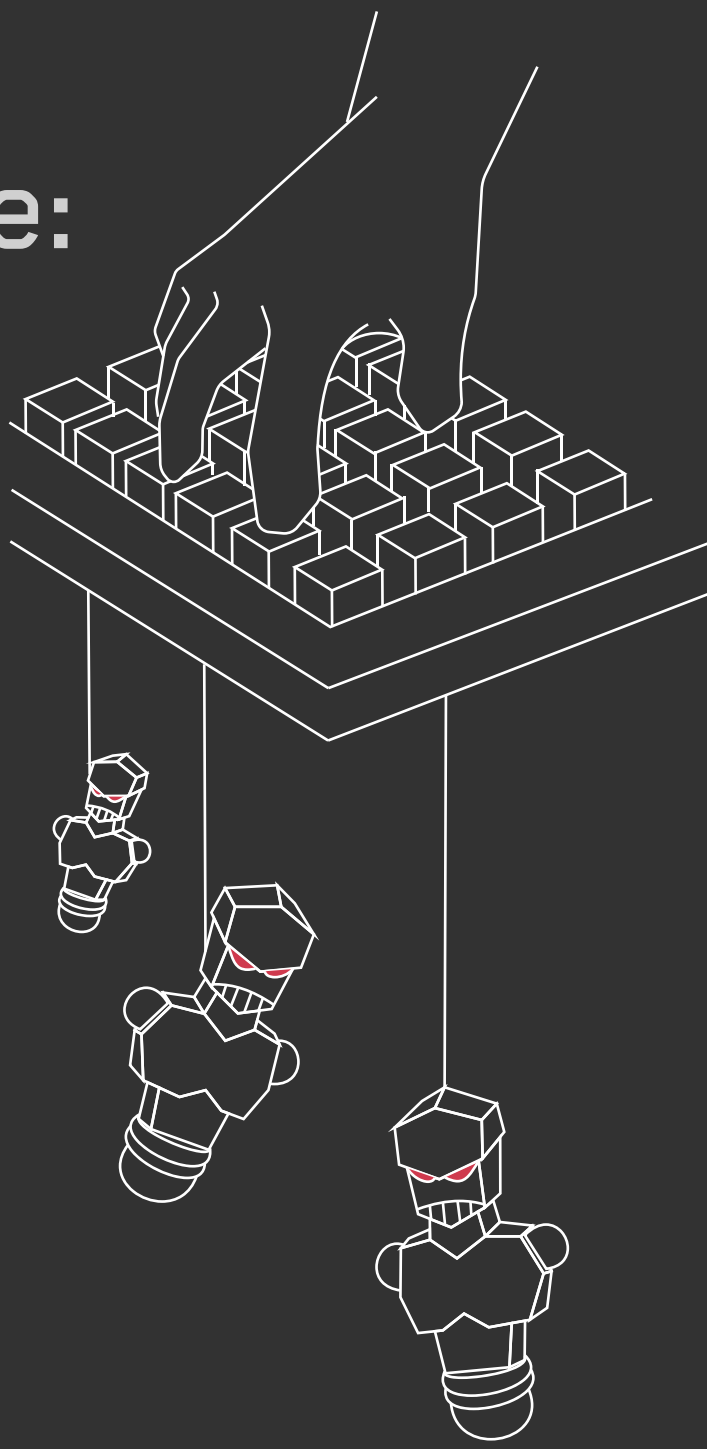

White Ops, Inc.

Association of National Advertisers

DECEMBER 2014

TABLE OF CONTENTS

05 About the Study

10 The Impact of Bots on Digital Media

18 Challenging Existing Assumptions: Bots Can End Up on Premium Sites

22 The Origin of Bots in the Media Supply Chain

28 How the Bots Blend In: Getting Targeted, Faking Metrics

33 How Bot Suppliers Get Away With It: Evasion

37 Sites That Only a Bot Could Love

41 When Publishers Are Victims Too: Ad Injection

44 Eliminating Bot Fraud: A Call to Action

Appendix A: Glossary of Terms

Appendix B: Constraints and Limitations

Appendix C: External Study Contributors

Appendix D: Illustrative Terms and Conditions

Special Thanks to the Following ANA Member Company Participants

ABOUT WHITE OPS

A pioneer in the detection of bots and malware on the web, White Ops develops new bot detection technologies to differentiate between bot and human interaction. Bot detection makes bot/human decisions in online advertising, publishing, enterprise business networks, e-commerce transactions, and financial systems. White Ops protects clients from bot fraud by cutting off sources of bad traffic to make bot and malware fraud unprofitable and unsustainable.

ABOUT THE ANA

The ANA (Association of National Advertisers) provides leadership that advances marketing excellence and shapes the future of the industry. Founded in 1910, the ANA's membership includes more than 640 companies with 10,000 brands that collectively spend over $250 billion in marketing and advertising. The ANA also includes the Business Marketing Association (BMA) and the Brand Activation Association (BAA), which operate as divisions of the ANA. The ANA advances the interests of marketers and promotes and protects the well-being of the marketing community.

ABOUT THE STUDY

EXECUTIVE SUMMARY

We expected to find bot-focused websites with nothing but a bot audience, but out of nearly three million websites covered in the study, mere thousands were completely built for bots. Most of the bots visited real websites run by real companies with real human visitors. Those bots inflated the monetized audiences at those sites by **5 to 50** percent.

Global advertisers will lose $6.3 billion to bots in 2015

At current bot rates, advertisers will lose approximately $6.3 billion globally to bots in 2015 (applying the bot levels observed across our study to the estimated $40 billion spent globally on display ads and the estimated $8.3 billion spent globally on video ads).

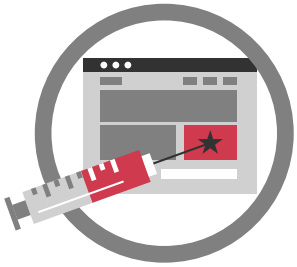

Ad fraud gets home users hacked

Bot traffic comes from everyday computers that have been hacked. Over **67 percent** of bot traffic observed in the study came from **residential** IP addresses. Bot traffickers remotely control home computers to generate ad fraud profits. Bots hijack browsers to masquerade as real users, blend in with human traffic, and generate more revenue.

Ad bots defeat user targeting

After infiltrating home computers with malware, cybercriminals make real money from their victims by installing ad bots. By using the computers of real people—people who are logged in to Gmail, sharing on Facebook, and buying on Amazon—the bots do not just blend in, they get targeted.

Bots coast on the credentials of the real users of the computers they hijack. Bots were observed to click more often (but not improbably more often) than real people. Sophisticated bots moved the mouse, making sure to move the cursor over ads. Bots put items in shopping carts and visited many sites to generate histories and cookies to appear more demographically appealing to advertisers and publishers.

BOTS ARE EVERYWHERE

...BUT NOT IN EQUAL NUMBERS

The study included a diverse range of brands, across nine vertical categories, with total annual U.S. ad budgets from under $10 million to over $1 billion, as measured by Kantar. The magnitude of the participants' ad spending had no correlation with the level of bots observed.

Bot percentages in our data skewed high:

At night

Approximately half the bots caught were not sophisticated enough to keep daylight hours.

In display

Bots accounted for 11 percent of all display impressions observed.

In video

Bots accounted for 23 percent of all video impressions observed.

In programmatic and retargeted inventory

Bot traffic in programmatic inventory averaged 17 percent. Bots consumed 19 percent of retargeted ads.

In sourced traffic

Third-party traffic sourcing resulted in 52 percent bot fraud.

In specific domain categories

Finance, family, and food domains showed increased bot traffic, ranging from 16 to 22 percent bots.

GOALS AND METHODOLOGY

Bots are software scripts in networks of computers that are controlled by a single entity as part of a botnet. The botnet controller can cause the computers in its botnet to execute a variety of behaviors and goals, including advertising fraud, online bank robbery, identity theft, and distributed denial of service (DDOS) attacks. When executing ad fraud, the botnet controller causes the computers in its botnet to render or click on ads, requiring advertisers to pay for a click-through or an ad impression that was never served to a real human.

Historically, huge volumes of ad fraud have been undetectable to advertisers. White Ops and the ANA worked with 36 ANA member organizations to analyze digital advertising campaign traffic over a period of 60 days between August 1 and September 30, 2014.

We used newly developed technologies that revealed bots and showed the true domain source of ad impressions. We studied **5.5 billion** impressions – the largest public study to date of bots in digital advertising.

White Ops provided guidance to all study participants, but contributors were permitted to select the type of ad traffic to be measured during the study. There was no uniform point of analysis, type, or percentage of traffic analyzed. Mandatory requirements were not placed upon participants. The sole unifying aspect of the methodology was the unique approach White Ops used to differentiate between a human and bot (machine-driven) request.

White Ops evaluated billions of impressions, discovered hundreds of millions of bots, and covered video and all types of display advertising. Display and video advertising purchased via direct, network, and programmatic channels were all evaluated.

This study examined traffic for 36 ANA participants from the following industry verticals: auto, beer/spirits, CPG, financial/insurance, hospitality, pharma, restaurant, retail, and technology.

36 Companies

181 Campaigns

3 Million Domains

5.5 Billion Impressions

60 Days

HOW WE USED THE DATA

This study is a baseline assessment of fraud in digital advertising, a threat that has emerged over the past decade.

We used new bot detection technology to collect data on ad fraud attacks. We compared data received from the participants to historical evidence from White Ops and external sources Chartbeat, Ghostery, and Grapeshot.

We aggregated evidence across different traffic types and analytic methods for the 36 participating ANA member organizations to minimize the influence of individual organizations in each of the samples.

The publicly announced study assessed premium digital advertising brands during a relatively slow portion of the advertising year, suggesting that the **bot measurements observed during this study underrepresent the overall level of bot fraud in the advertising ecosystem.**

In this report, we provide recommendations to assist advertisers, agencies, and publishers in developing defenses against the increasing threat of digital ad fraud.

THE IMPACT OF BOTS ON DIGITAL MEDIA

BOTS WERE IN EVERY KIND OF CAMPAIGN WE STUDIED

Bots were not deterred by the technical challenges of consuming video inventory. Video CPMs are typically much higher than display CPMs, and video inventory contained over **twice** the percentage of bot fraud.

Measurement of bots in retargeting campaigns and bot engagement metrics indicate that higher CPM inventory may concentrate populations of sophisticated bots. The most elite botnet operators appear to customize their bots to capitalize on the greater dollar opportunity available in higher CPM inventory (see *Bots Faked All of the Engagement and Viewability Metrics That We Measured*, page 29).

OVERALL STUDY BOT PERCENTAGES

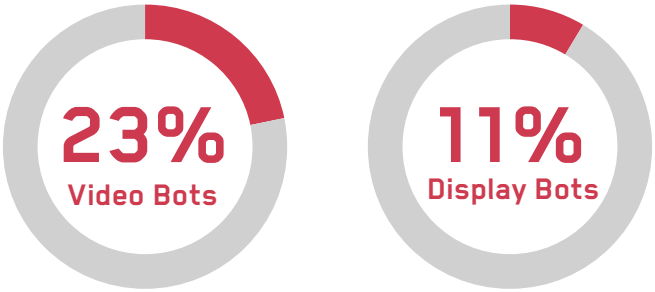

CASE STUDY

Participant Hit by High Bot Levels in Premium Display and Video Ads

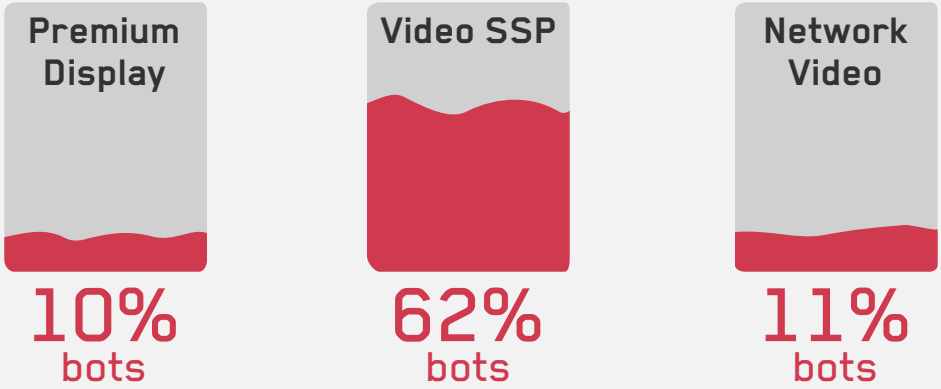

The agency for one of the CPG participants ran 12 placements on sites owned and operated by a major U.S. cable and media concern, which had 10 percent bots. This same CPG participant bought video advertising on both a publicly traded video supply-side platform (SSP) and a leading Internet portal, with bot levels of 62 percent and 11 percent, respectively.

BOT FRAUD VARIES ACROSS CAMPAIGNS AND PLACEMENTS

Because bot percentages vary unpredictably across campaigns and placements, identifying and comparing the real cost of the bot fraud problem is complex. One successful strategy is to examine fraud losses at the placement level rather than looking at total losses within a brand or company-wide fraud loss.

RECOMMENDATION

Troubleshoot placements that lose more ad spend than others.

Using granular information (e.g., "Placements I, K, L, and M have below-average bots, but we are losing 33 percent of impressions on Placement D to bots"), advertisers can quickly diagnose and eliminate large monetary losses to bots even if the brand's overall level of bot traffic is low.

Cost-per-human (CPH) measures the actual cost of human impressions after accounting for loss due to bot fraud. Advertisers can compare CPH values across placements, campaigns, brands, and traffic sources to understand and communicate the real cost of reaching a human audience goal.

FINDING

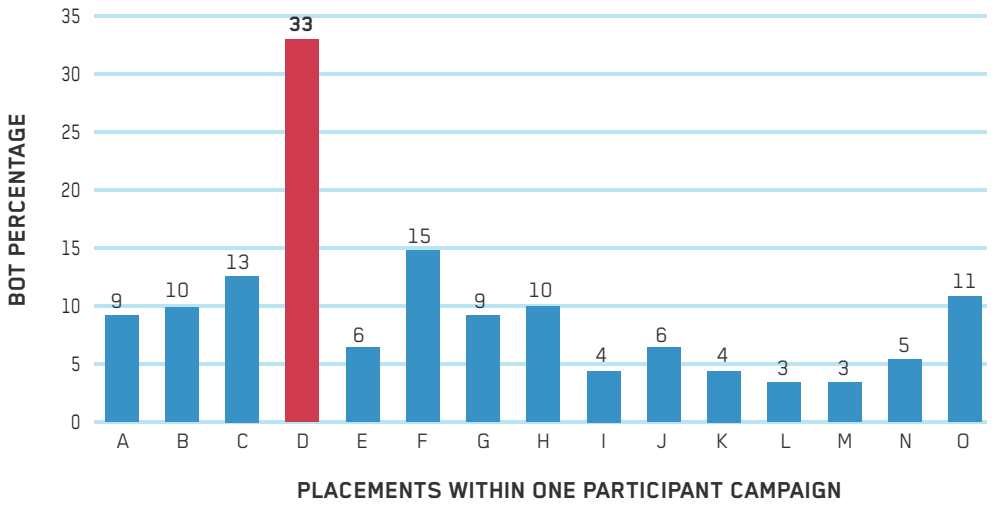

Figure 1: Placement Bot Percentage Can Vary Within a Single Brand

CPH
Cost per Human

Measures the actual cost of each thousand human impressions after accounting for loss due to bot traffic

BOT FRAUD VARIES OVER TIME

Not all bot operations are created equal. Half of bots are not sophisticated enough to keep normal waking hours.

FINDING

At night, bot percentages were higher.

Total bots observed were higher during the day but were a lower proportion of overall traffic.

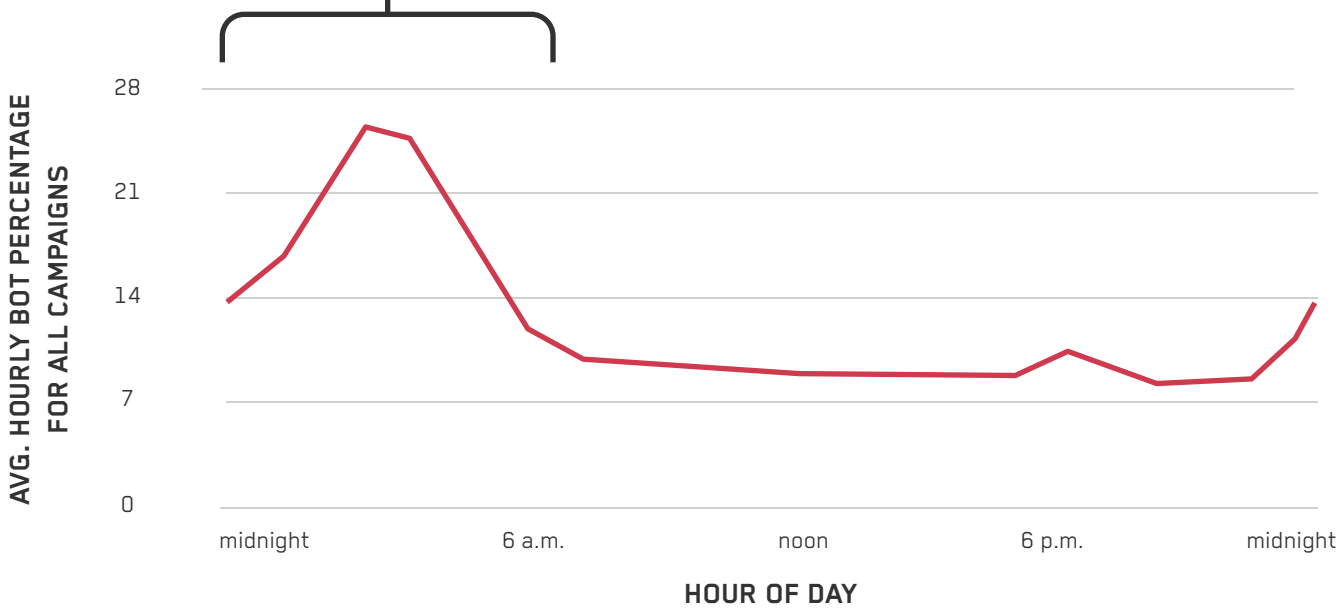

Figure 2: Bot Traffic Time-of-Day Pattern Across Study Time Period

Local time determination was made using IP geolocation data provided by Maxmind.

RECOMMENDATION

Consider day-parting to reduce time-of-day bot percentage spikes.

EFFORTS TO INCREASE REACH ALSO INCREASE BOTS

In online advertising, reach refers to the number of unique individuals exposed to an ad. When study participants attempted to increase their reach with run-of-network (RON) campaigns, those campaigns were exposed to higher than average bot fraud, averaging 16 percent.

When advertisers demand more traffic, often at the end of the week, month, and quarter, the differential between available humans and advertiser demand for traffic can be made up with bots.

Bots can make it look easy to reach high volumes of specific audiences. A bot can look like a sports fan, someone with a six-figure income, someone interested in buying a car, or a grandparent looking for holiday gifts for grandchildren.

White Ops has historically observed that campaigns around time-sensitive releases such as retail sales, movies, and TV shows are unusually vulnerable to bot activity because they have very specific delivery windows that can exacerbate the bot problem.

CASE STUDY

Bots Spiked Every Saturday

For one beer/spirits participant, a bot spike occurred at the end of every week of a campaign, with bots spiking at noon (PST) on Saturday, uniformly increasing from zero to 800 bots per hour before dropping back to zero bots per hour after the peak at noon on Saturday.

Bot levels dropped to nearly zero for all other days of the week (shown in the nearly empty columns of the graph, at right). The bot spikes on Saturdays comprised 95 percent of all the bot fraud for this campaign.

- Each dot placed on the y-axis represents the number of bots detected during each hour of the study.

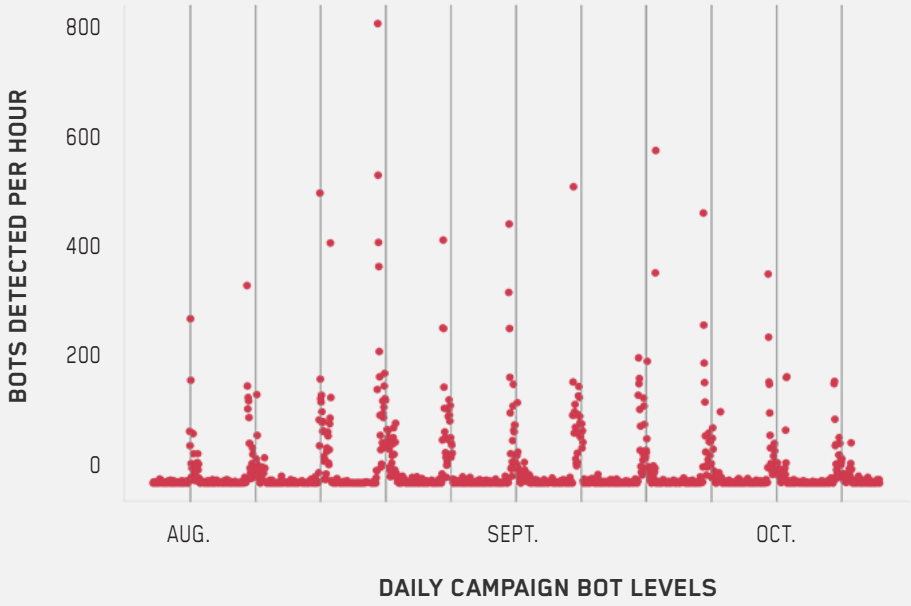

Figure 3: Bots Spiked in This Campaign Every Saturday at 12 P.M.

BOT TRAFFIC VARIES BY DOMAIN CATEGORY

No consistent variation in bot traffic percentage was evident across participant industry verticals. However, when we correlated our bot results with domain content analysis from Grapeshot, bot traffic varied notably across domain categories where the ads were served.

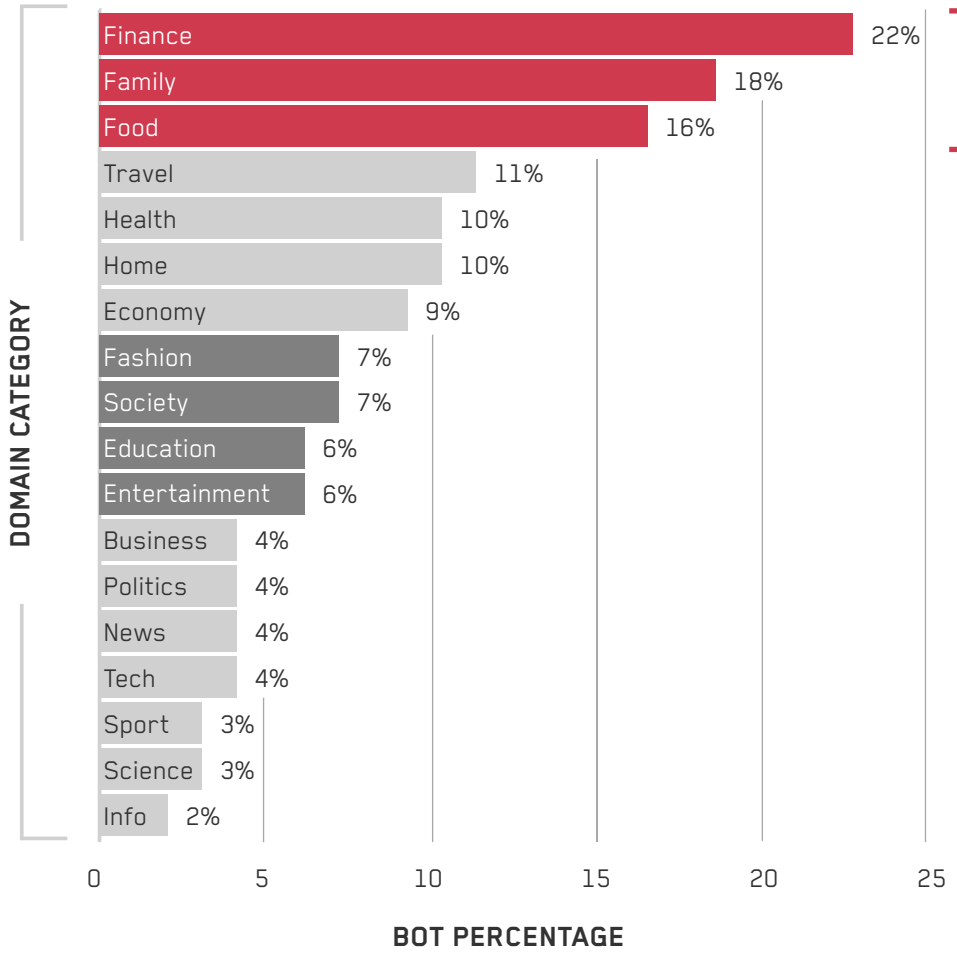

Finance, family, and food domains had the highest percentages of bots, with 16–22 percent bots. Tech, sport, and science had the lowest bot percentages, ranging from 3–4 percent bots.

Figure 4: Finance, Family, and Food Domains Had Higher Bot Percentages

Domain category data was provided by Grapeshot.

AD FRAUD TAKES A BIG BITE OUT OF VIDEO CAMPAIGNS

FINDING

Bots accounted for 23 percent of all video impressions observed.

Bot traffic ranged from 2 percent to 100 percent in video placements. Bot traffic by participant campaign (sometimes consisting of multiple placements) was as high as 63 percent. Some participants ran multiple video campaigns, with up to 90 million video ad impressions from a single participant. There was no correlation found between campaign volume and bot levels.

Video ads were vulnerable to non-bot-driven ad fraud in addition to bot traffic. **Some of the highest impression-volume video ad campaigns took huge hits from bot traffic as well as from video autoplay adware** (see *Adware: Not All Fraud Is Robotic,* page 24). Instead of deploying bots, the video autoplay adware used humans who had minimal or no control of the adware applications to fraudulently consume advertising inventory.

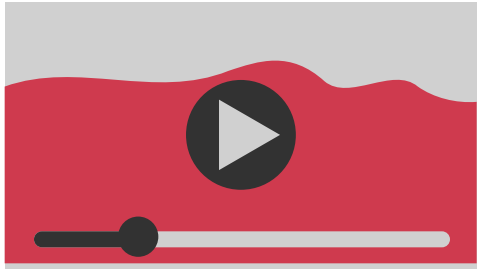

Up to **63%**

Average Bot Traffic in Video Ad Campaigns by Participant

RECOMMENDATION

Without the ad fraud, digital video advertising holds great promise. Agencies and advertisers can capture the immense potential value of online video advertising campaigns to meet marketing goals by putting in place quality control mechanisms.

These same quality assurance measures will help protect programmatic display campaigns from ad fraud.

To ensure the integrity of video and programmatic display campaigns:

- Implement continuous fraud monitoring.
- Use bot detection to ensure that sites are not sourcing traffic.
- Monitor for all types of ad fraud, including adware and bot traffic.

SOURCING TRAFFIC INCREASES BOT LEVELS

Publishers sometimes use traffic sourcing (any method by which publishers acquire more visitors through third parties) to improve measured audience levels at their sites.

We compared indicators of traffic sourcing to study-wide bot percentages in traffic. **Indicators of traffic sourcing predicted high bot percentages more frequently than almost any other behavioral factor.**

The high bot percentage of sourced traffic remained stable over the length of the study. Bot activity on specific sources ranged from fully human (0 percent bots) to 100 percent bots. This range included disclosed incentivized traffic.

Well-known publishers and premium publishers were not immune to high bot levels in sourced traffic.

Sourced Traffic Averaged 52% Bots

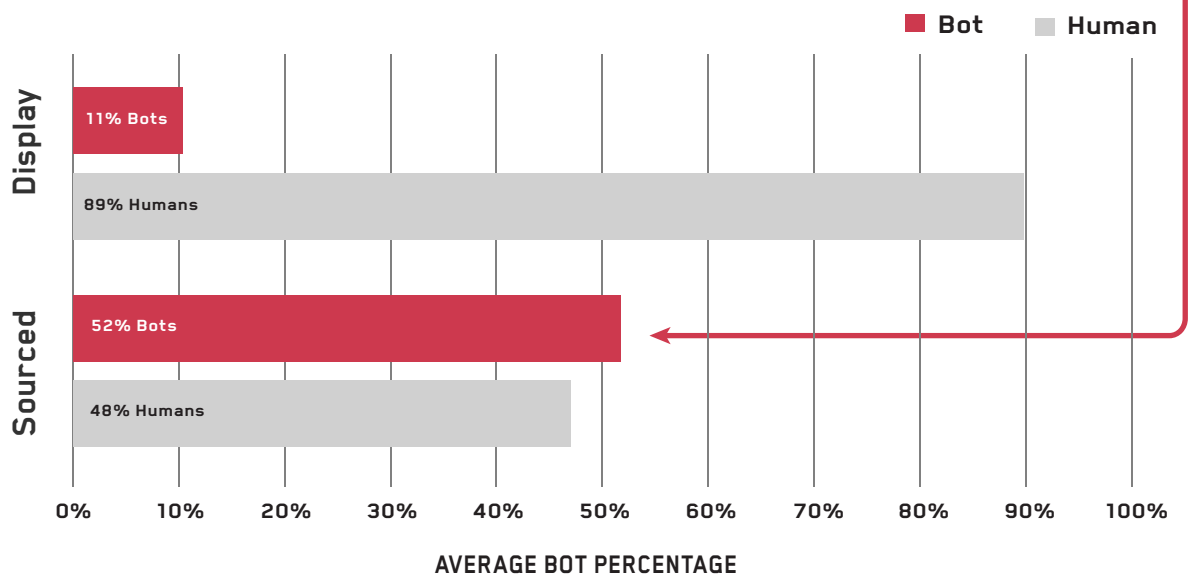

Figure 5: Traffic Sourcing Generated More Bot Traffic Than Human Traffic

CHALLENGING EXISTING ASSUMPTIONS: BOTS CAN END UP ON PREMIUM SITES

BOTS GET INTO PREMIUM BUYS

Significant bot levels affected all tiers and types of publishers. Premium, direct-buy display advertising campaigns at many well-known domains showed bot percentages higher than 10 percent.

Ad fraud in premium publisher traffic came from bot traffic and ad injection (the unauthorized placing of ads on sites where they do not belong).

Advertisers who assume that traffic to premium publishers is free of bots risk losing large amounts to intentional or unintentional bot fraud.

CASE STUDY

Premium Publisher Serves Up 19% Bots

A CPG participant purchased 230,000 impressions from a premium U.S. media company. Traffic from the site averaged 19 percent bots.

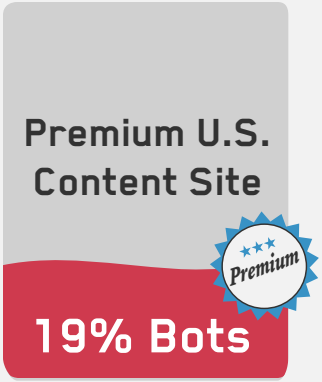

CASE STUDY

Direct Buy at Premium Publisher Yielded 98% Bots in Video Ad Campaign

One premium, well-known publisher in the lifestyle industry vertical employed a web page layout consisting of a single large video player at the top of the page. Seemingly random selections of content surrounded the autoplaying video on the page.

On this publisher page, video ads for an auto participant in the study were consumed by a 98 percent bot audience. Out of almost 4,000 total video impressions from the placement, fewer than 100 were served to humans.

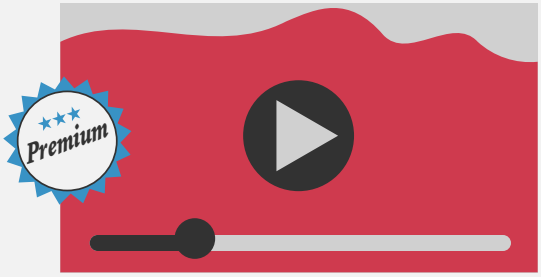

98% Bots

THE GAME OF AVERAGES: MIXING CLEAN PLACEMENTS WITH BOT-HEAVY PLACEMENTS LOWERS THE OVERALL AVERAGE

CASE STUDY

16 to 64% Bot Traffic Hit Direct Buy, Premium Campaigns

An agency for a participant in the retail vertical placed a direct buy on sites owned and operated by a well-known U.S. media company. Sixty placements contained an average of 17 percent bots. About half of the placements were very clean; the other half ranged from 16 to 64 percent bots.

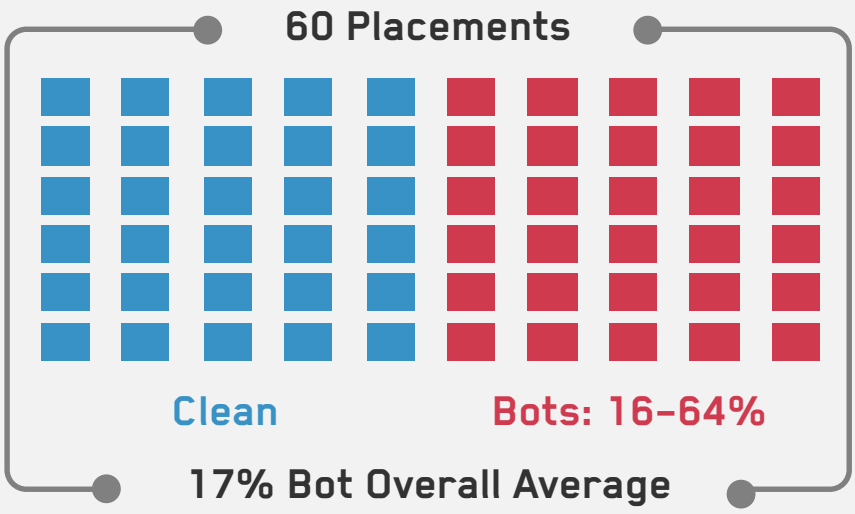

TO REDUCE BOTS IN PREMIUM CAMPAIGNS, ADVERTISERS MUST UPDATE ASSUMPTIONS

The quality of some premium, well-known publishers may have degraded since they established their reputations in the early days of the web. The reputation of the publisher is no longer a reliable benchmark to predict bot traffic levels.

RECOMMENDATION

Use technology to validate all assumptions.

- To avoid paying for bots in premium campaigns, use third-party fraud detection to validate or disprove assumptions about ad buys from all sources, including premium or Tier 1 publishers and trusted sources.
- Avoid making assumptions about traffic quality based on a publisher's premium or tier classification.

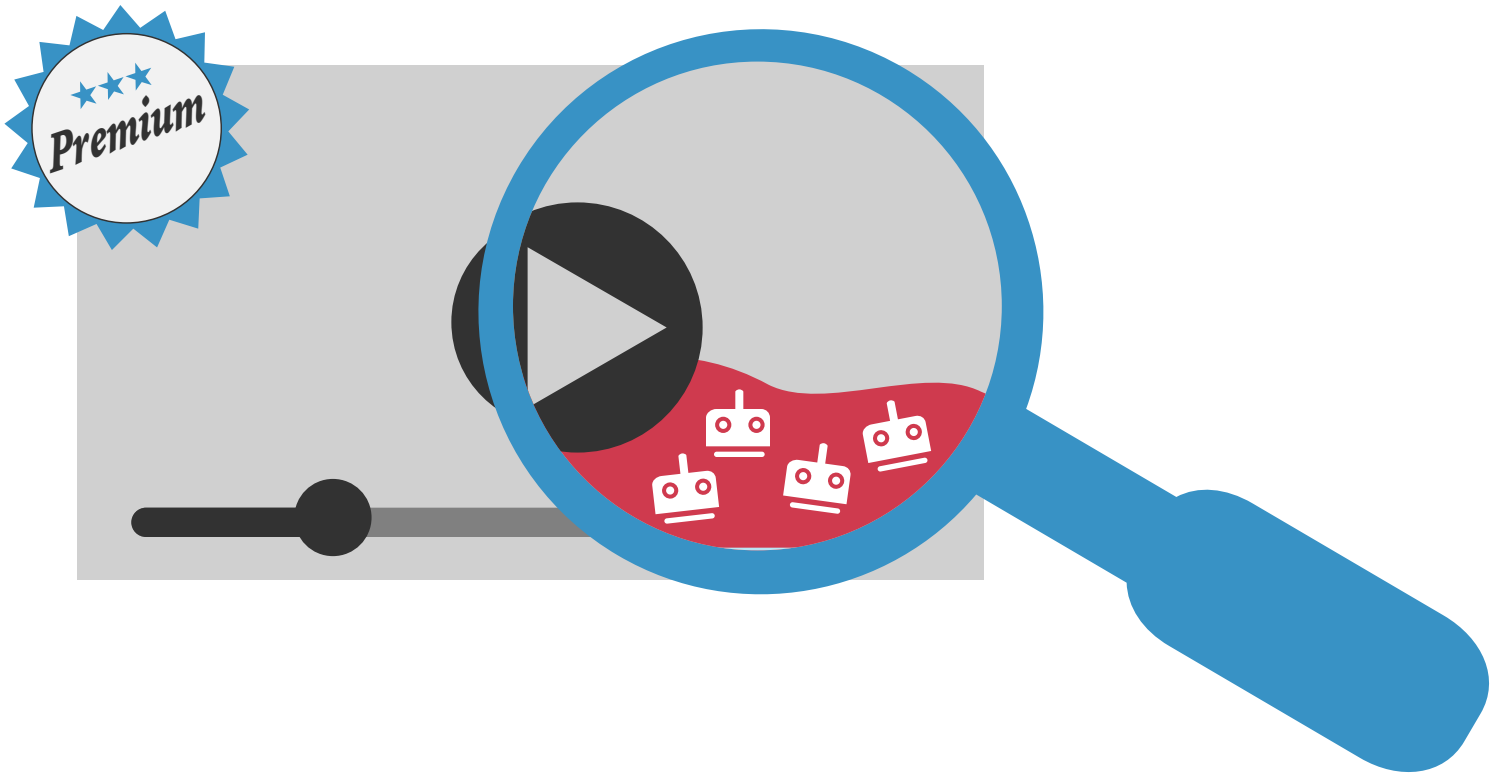

THE ORIGIN OF BOTS IN THE MEDIA SUPPLY CHAIN

BOTS ARE BUILT TO FOOL ALL STAKEHOLDERS IN THE DIGITAL ADVERTISING SUPPLY CHAIN

Some links in the bot supply chain are unaware of the bots in their traffic and do not intend to profit illicitly, while others actively encourage and concentrate bot traffic to increase profits.

BOT SUPPLIERS FOOL ALL STAKEHOLDERS IN THE DIGITAL ADVERTISING CHAIN:

- Advertiser's marketing team
- Agency's media planning team
- Agency's media buying team
- Agency's ad server
- Agency's analytics and optimization team
- Publisher's ad operations team
- Publisher's sales organization
- Publisher's ad server
- Third-party ad verification services
- Advertiser's internal audit controls

Each bot heist occurs in a matter of milliseconds and is often undetectable to the victim. Bot suppliers repeat these automated heists at the scale of hundreds of millions of times per day, creating enormous scope and reach in their fraud.

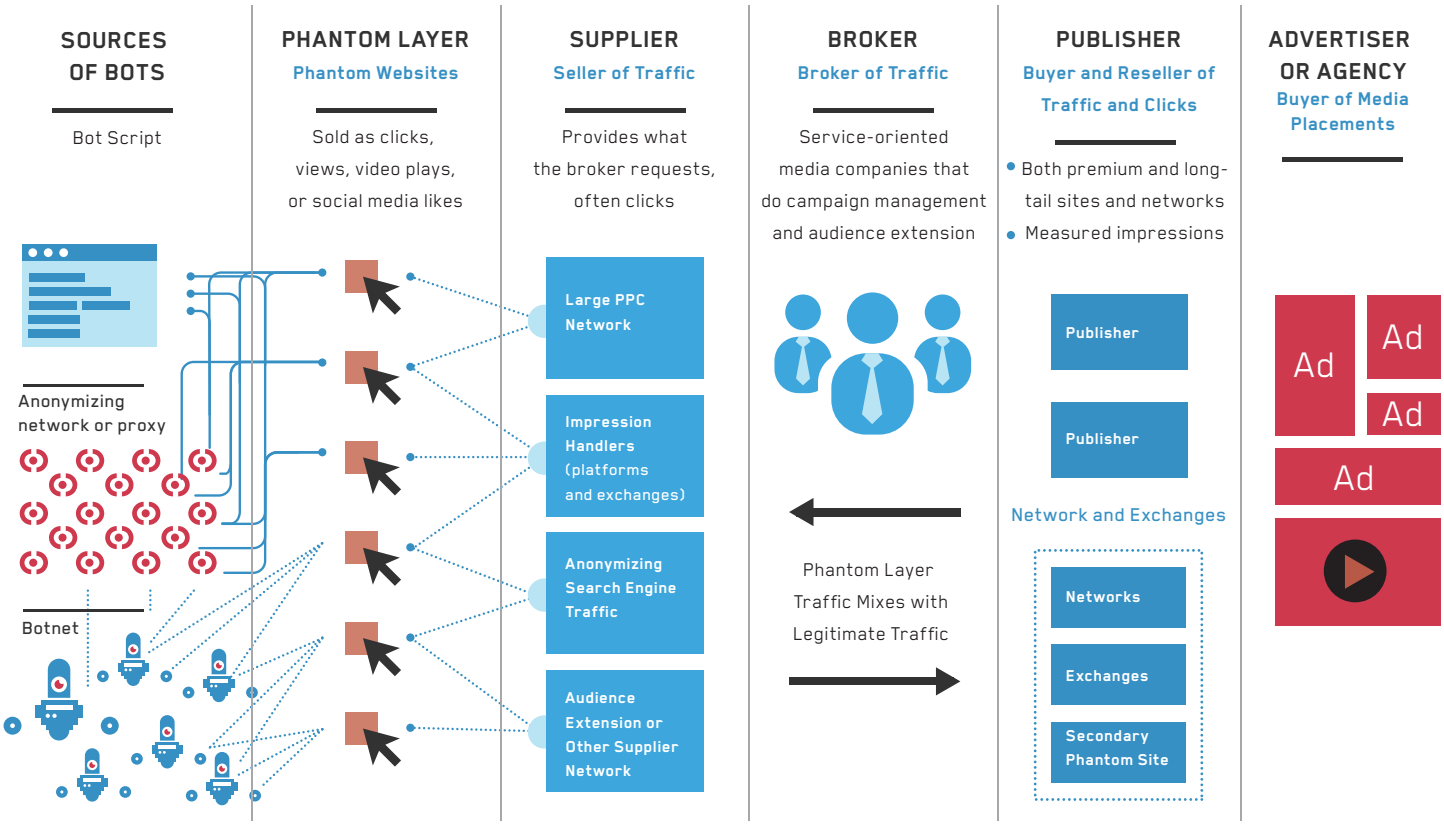

Figure 6: A Phantom Layer Generates and Launders Fraudulent Web Traffic

Bot impressions originate from malicious bot suppliers and pass through both legitimate and phantom layer elements of the digital advertising ecosystem. Phantom layer elements are websites operated specifically for the purposes of laundering ad fraud.

ADWARE: NOT ALL AD FRAUD IS ROBOTIC

Adware is software, often automatically installed on user devices, that serves visible or hidden ads to users to boost ad consumption.

For several study participants, a significant amount of adware-enabled activity that occurred was classified as ad fraud, not bot fraud.

The adware traffickers did not use bots to create the undesirable ad impressions. However, the adware activity was unsanctioned, fraudulent, and harmful to everyday computer users and advertisers. The adware represented enormous revenue to the publishers for impressions that were undesirable to the advertisers.

The adware that affected these participants was very similar to bots. The main difference was that the adware created a pop-under window visible to the user until the user closed the pop-under, at which point the adware continued to operate in the background without the user's knowledge.

CASE STUDY

Adware Preys on Everyday Computer Users and Consumes Digital Media Budgets

One participant's video ad campaign was delivered via 10 million adware impressions from a single adware trafficker within the first week of the study. Of the campaign's nearly 90 million impressions, 7 percent were natural human impressions, and 93 percent were fraud.

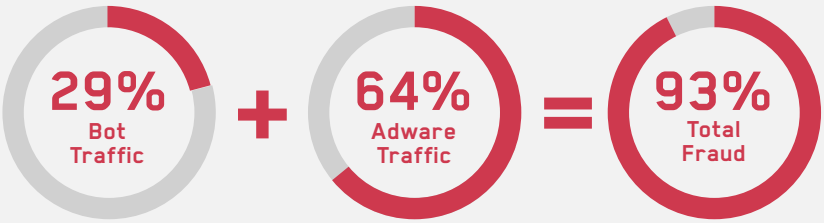

The publisher of the adware provided a video ad-supported service that required the user to download adware software and, in some cases, was seen to pay for the non-consensual installation of the adware software on unsuspecting users' computers. The adware ran ads continuously in a browser in the background of the user's computer, one video at a time. In most cases, the ads were entirely hidden from the user.

The adware initially showed a pop-under to the user that played video ads. The adware changed the volume for itself to zero while playing audio, leaving volume controls for other software untouched. After the user closed the pop-up, the adware continued playing ads with silenced audio. After restart and login of the user's computer, the adware software autoplayed video ads regardless of whether the user reopened the adware site or application. The adware's autoplay functionality was unsanctioned and uncontrollable by the user.

ADWARE: NOT ALL ATTACKS ARE THE SAME

CASE STUDY

Adware Attacks Vary in Severity

A second adware site autoplayed video ads without sound, but only when the user could see the ads. This adware did not continue playing ads to completion when the user closed the ad pop-up, allowing the user more control. In this case, of a CPG participant's seven million impressions, 59 percent were natural human impressions and 41 percent were fraud.

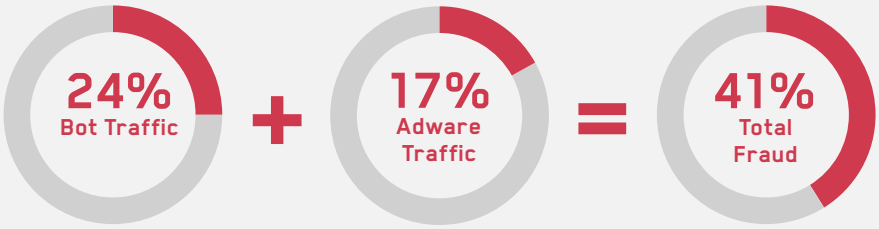

TRAFFIC BREAKDOWN FOR ONE PARTICIPANT'S VIDEO CAMPAIGNS

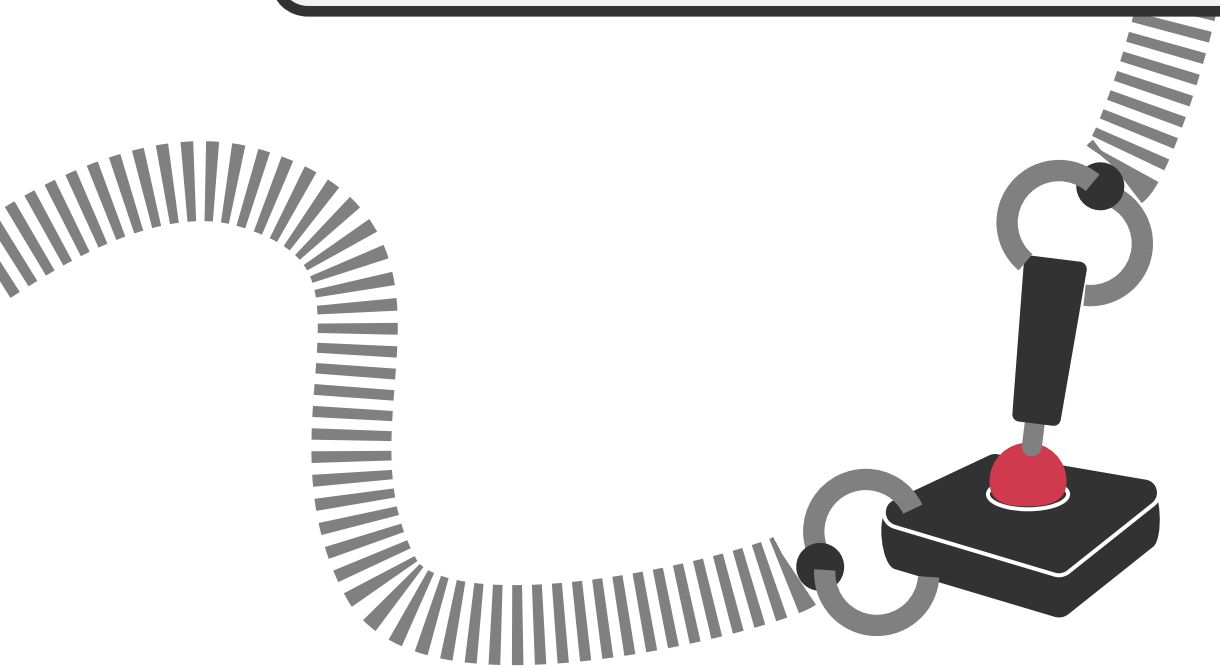

MOST BOTS COME FROM RESIDENTIAL IPs

Some advertisers attempt to reduce the fraud problem using IP blacklists. Over time, bot suppliers evolve their bots against the metrics that were used to create and defeat the blacklist. Domain blacklists, geographic blacklists, and browser or retargeting lists we studied were not effective at stopping the majority of bot fraud.

Bot Source by IP Type

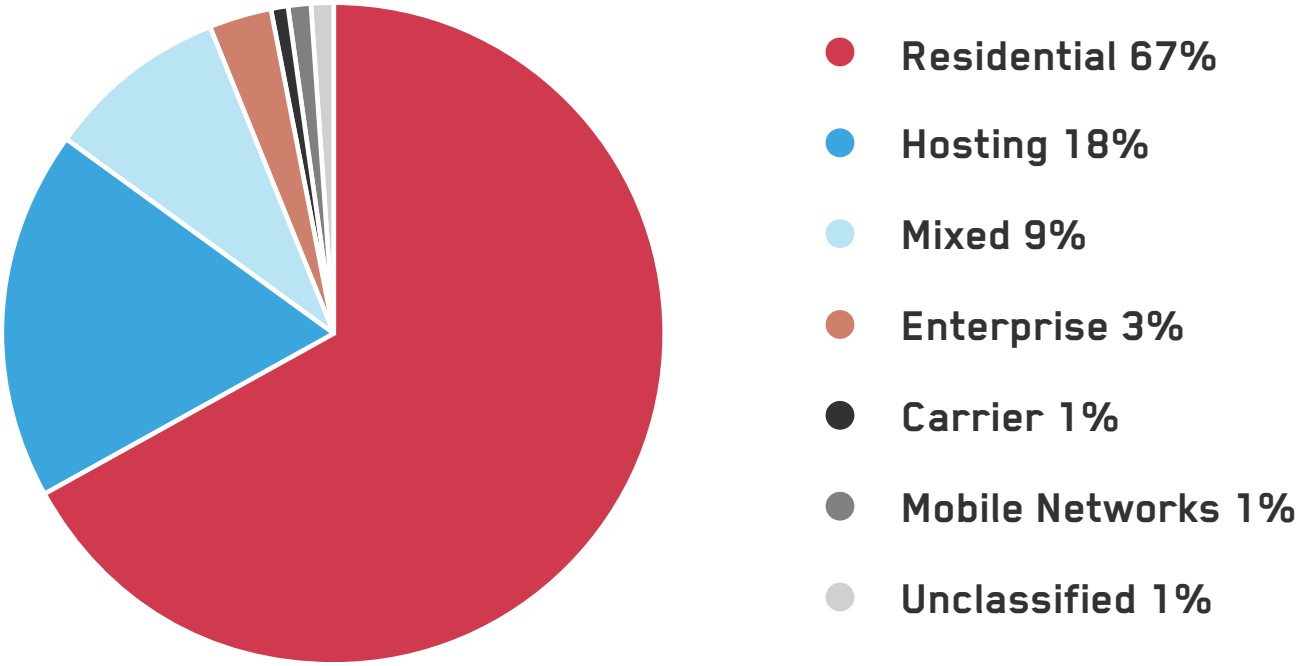

Figure 7: Botnet Operators Break Into Everyday Users' Computers to Remotely Drive Bot Fraud

One of the earliest resources created to address the bot problem was the IAB/ABC International Spiders and Bots List. This list was designed to detect non-human traffic and prevent such traffic from being counted in web analytics and is still in existence today. It enables filtering of non-human activity that can significantly inflate ad impression and site traffic counts and is updated monthly. The end result is a more transparent and accurate measurement for ad impressions and site traffic claims. However, the IAB/ABC International Spiders and Bots List was not designed to track criminal botnets.

BOTNET CONTROLLERS HIJACK EVERYDAY USERS´ IDENTITIES AND MACHINES

Bot traffic could be coming from the computer you are using right now. Over 67 percent of the bot traffic observed in the study came from residential IP addresses.

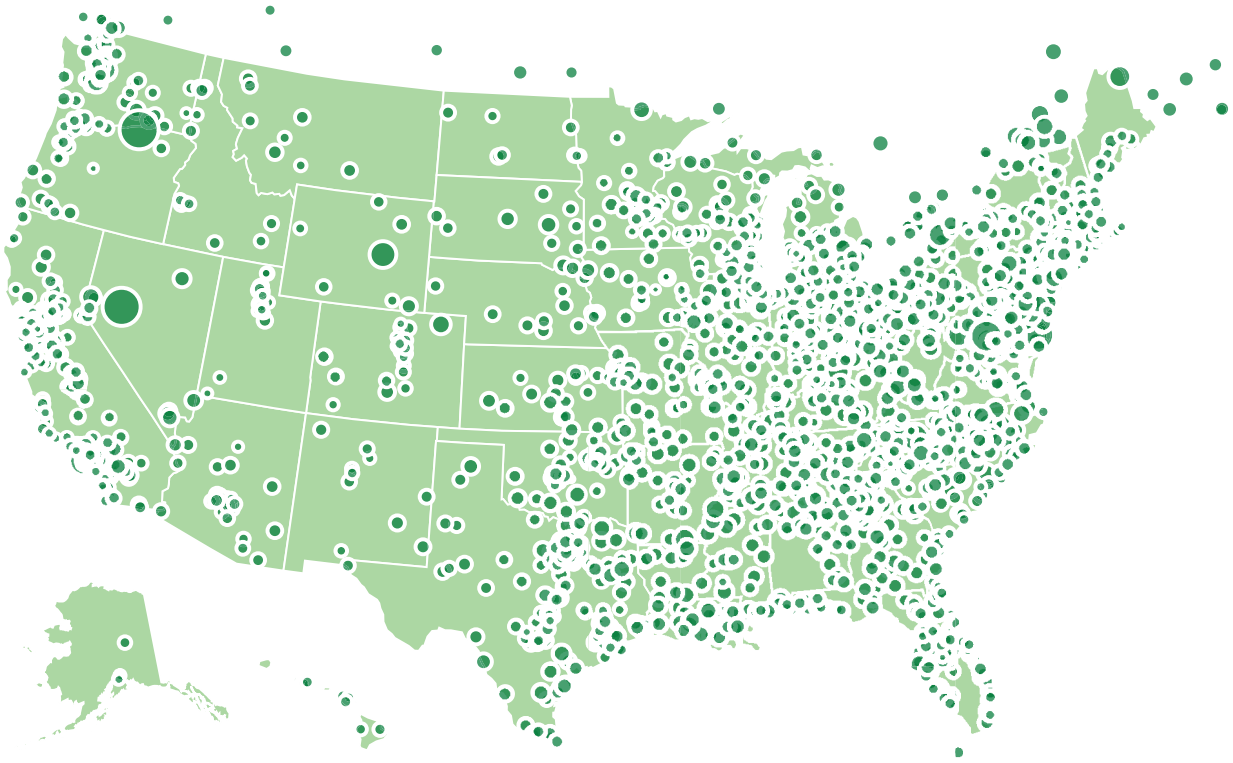

Figure 8: Residential IP Address Sources of Bot Fraud Are Distributed Throughout the U.S.

Bot traffickers hack U.S. home computers to get access to U.S. IP addresses and cookies. A small percentage of highly compromised computers create the bulk of the bot traffic.

HOW THE BOTS BLEND IN: GETTING TARGETED, FAKING METRICS

BOTS FAKED ALL OF THE ENGAGEMENT AND VIEWABILITY METRICS THAT WE MEASURED

In partnership with Chartbeat, White Ops compared engagement metrics between bots and humans.

High-bot sites had more abundant, less engaged bots; bots were more engaged than humans on low-bot sites

Bots on low-bot sites were sparse but highly engaged. On these sites, bots stayed engaged on the page 5 percent longer than humans and scrolled down in the page 12 percent less than humans. On high-bot sites, bots remained engaged on a page only 14 percent as long as the average human.

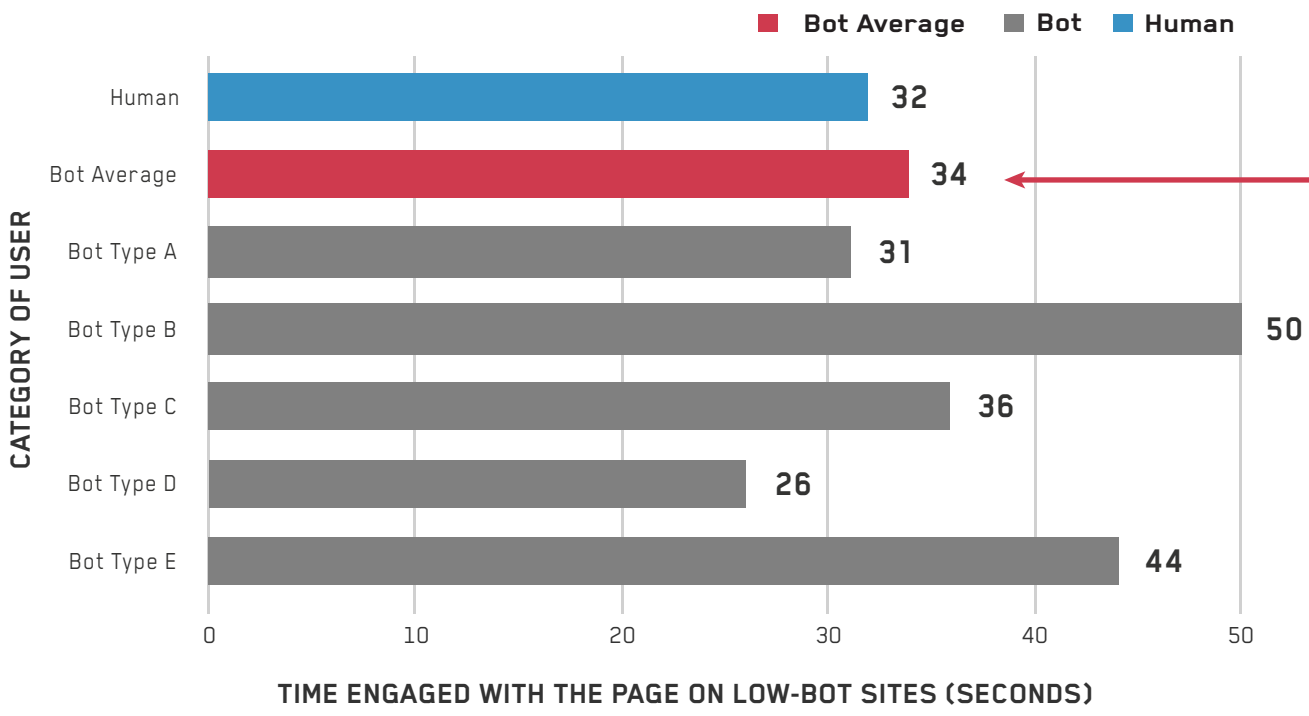

Figure 9: Bots Are Often More Engaged Than Humans

Engagement measurements were provided by Chartbeat.

VIEWABILITY DOES NOT ENSURE HUMANITY

Bots are built to run a web browser. When the bot and browser consume media, the browser reports that it is viewable even when it is not actually being rendered to the screen. This innate characteristic of bots means that **bots show high viewability by default.**

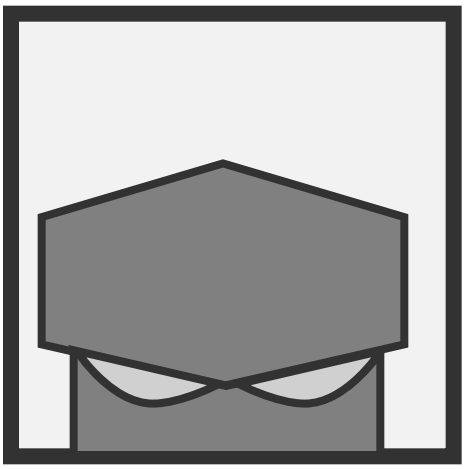

The five most common bot types consumed more ads than humans

Working with Chartbeat, we compared bot and human viewability at 87 high-humanity sites.

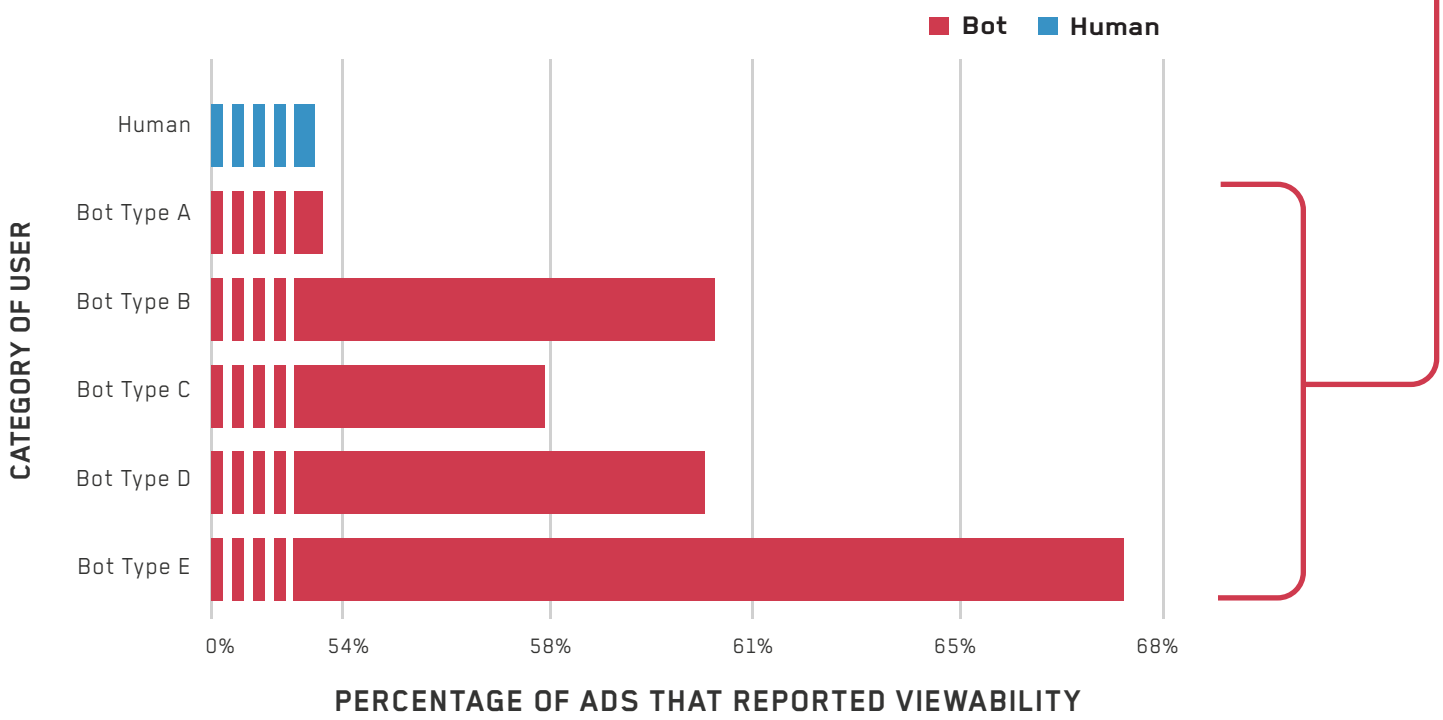

Figure 10: Bots Show High Viewability by Default

Viewability measurements were provided by Chartbeat.

BOTS ARE GETTING TARGETED AND RETARGETED

A bot typically visits more websites and consumes more ads than a human. When bots overshoot or miss the mark on engagement metrics, they get caught. When bots match the advertiser's targeted human engagement metrics, they can avoid detection and increase revenue.

Bots consumed 19 percent of retargeted ads in the study. Retargeting is the process of delivering ads to particular users based on their previous online activity.

Web browsers encode tracking data called cookies that allow advertisers and publishers to track and remember their online visitors. Bot browsers generate their own cookies or build cookie profiles designed to make them look more human and more tempting to advertisers. Bots emulating engaged human behavior have high-value cookies and get targeted by advertisers, amplifying botnet revenue potential.

Bots can get re-cookied without even trying. In certain advanced botnets, bots access and use browser cookies generated by the computer's actual human user, making it even easier for the bot to get picked up by a retargeting campaign.

In many cases, re-cookied bots are added to granular consumer segmentation lists, creating a cycle of baking bots into many ad-tech platform audience models.

CASE STUDY

Retargeting Bot Prevalence Was Much Higher Than Overall Bot Percentage

Overall campaign traffic for one participant contained 17 percent bots. In contrast, the participant's retargeted campaigns were composed of 55 percent bots.

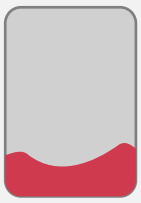

17%
Bots
Participant's Overall Traffic

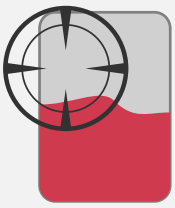

55%
Bots
Participant's Retargeted Campaigns

Similarly, campaign traffic for another participant contained 19 percent bots. This participant's retargeted campaigns were composed of 27 percent bots.

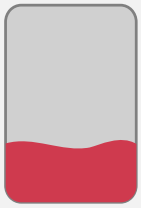

19%
Bots
Participant's Overall Traffic

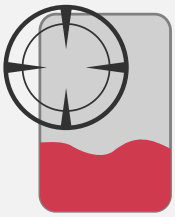

27%
Bots
Participant's Retargeted Campaigns

PROGRAMMATIC INVENTORY IS VULNERABLE TO BOT TRAFFIC EVEN WHEN OBTAINED FROM TRUSTED SOURCES

FINDING

Bot percentage varied widely when served from known programmatic ad server domains, with no identifiable predictor of bot traffic levels. For 18 of the 36 study participants, three well-known programmatic ad exchanges supplied programmatic traffic with over 90 percent bots (see case study at right).

The study average bot percentage for programmatic placements was 17 percent.

Average programmatic bot traffic for participants ranged from 3 percent to 31 percent.

RECOMMENDATION

Reputation and trust levels cannot predict the percentage of bot traffic a supplier will have. Buyers can take action to protect programmatic buys:

CASE STUDY

One Publisher Funneled Over 90% Bot Traffic Through DSPs to Half of Study Participants

A bot site used the opacity of programmatic display traffic sourcing through demand side platforms (DSPs) to systematically defraud advertisers.

Agencies for 18 of the 36 study participants bought inventory through three well-known, trusted DSPs, with each agency targeting different consumer profiles and segments. Ads from all 18 of these participants ran on a common publisher and showed consistent bot percentages greater than 90 percent.

Where's the Content?

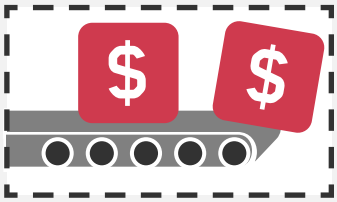

Programmatic Exchange

Ads at this site yielded payments to the publisher but were not consumed by humans. The bot supplier designed this site to efficiently capture advertiser dollars, placing ads in rows down the length of the web page, with no actual content on the page.

The site was designed to maximize fraudulent gains while minimizing botnet resource requirements.

Continuously monitor and troubleshoot programmatic buys.

- Monitor all ad inventory, even from trusted sources.
- Require sources to monitor traffic they provide.
- Pause or troubleshoot campaigns that do not meet percent-humanity targets.

HOW BOT SUPPLIERS GET AWAY WITH IT: EVASION

BOT SUPPLIERS EVADE STUDY

White Ops believes the overall bot levels identified in this study were depressed due to the public announcement of the study and general industry awareness preceding it.

Entities throughout the bot supply chain can react when alerted to traffic auditing and monitoring. The good news is that just the act of monitoring campaigns can suppress some bot activity.

CASE STUDY

Bot Traffic Dodges Public Study

With participants' acknowledgment, we continued bot detection through the month of September in a covert study phase.

For one participant, the overall bot percentage was 41 percent on August 2, the second day of the publicly announced study period. This number dropped dramatically two days later, to 4 percent on August 4. The participant's bot percentage remained low until the end of August, the publicly announced end of the study.

The study continued covertly throughout the month of September. By September 9, the participant's bot percentage climbed back to 38 percent and remained far higher than the study average bot percentage through the end of the month.

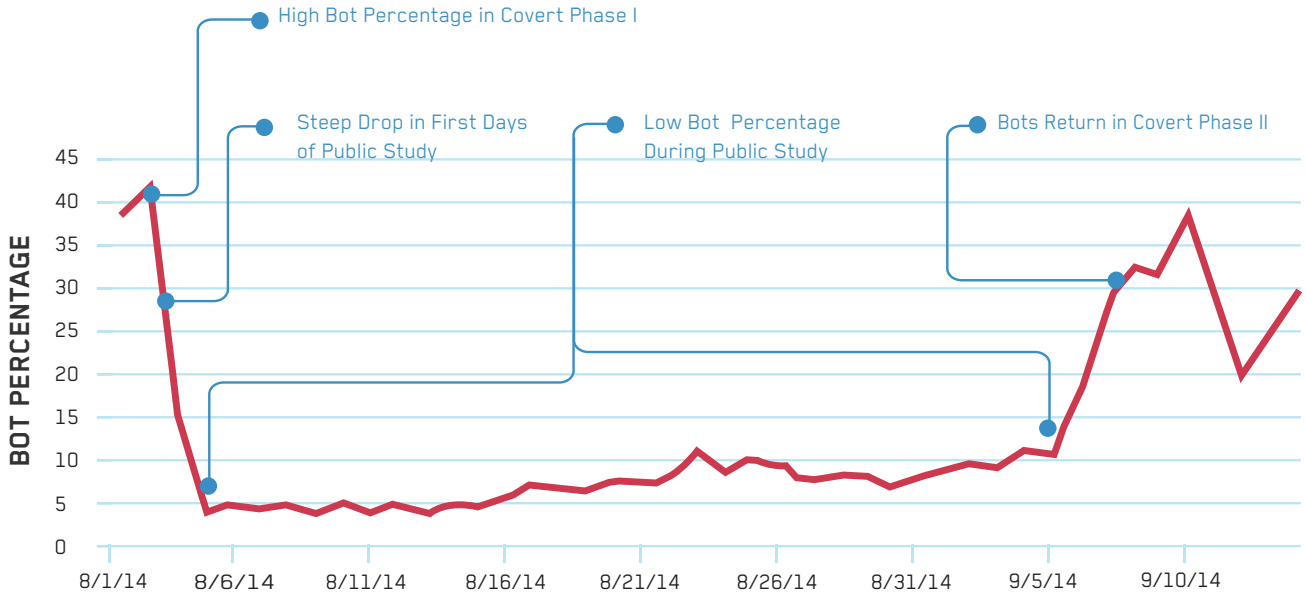

Figure 11: Bot Traffic Supply for This Participant Dodged Our Study

BOT SUPPLIERS ACTIVELY DISGUISE BOT TRAFFIC DURING CAMPAIGN MONITORING

The incentivized human type of ad impression does not usually represent a major loss of advertising dollars by itself because it does not scale as cheaply or as readily as bot traffic. However, in one case, a supplier leveraged this source of counterfeit impressions in an attempt to conceal bot traffic levels during a brand's audit.

CASE STUDY

Bot-Cloaking Mechanics

Figure 12: Bot Trafficker Actively Evaded Advertiser's Audit

Evasive maneuvering by bot suppliers also occurred during a campaign audit independent of this study. After initiating campaign monitoring, a brand informed its traffic supplier that it was aware of bot percentages above 90 percent in the supplier's traffic. Subsequently, the supplier's traffic to that single brand dropped to just 4 percent bots.

The bot supplier's traffic remained higher than 90 percent bots for other buyers, including traffic for one study participant that was concurrently running a campaign (bought by its ad agency on a well-known video SSP) using traffic sourced from the same bot supplier.

Fraud percentage in the supplier's traffic to the brand did not actually decrease to 4 percent. The supplier routed incentivized traffic (human traffic that is paid to view or click on ads) to the brand that had raised awareness of the original bot activity.

After being informed of the high bot traffic levels, the bot supplier executed a complex series of audit-defeating measures to make its traffic appear to be legitimate:

- Switching arbitrarily between real human traffic and bots
- Responding to complaints about traffic bot levels
- Maintaining similar traffic volumes for both bot and incentivized human traffic

TO DETER AND DETECT AD FRAUD, ADVERTISERS MUST BOTH CONSPICUOUSLY AND COVERTLY MONITOR FOR FRAUD

Bot traffic percentages sometimes drop when bot suppliers become aware of scrutiny. Advertisers can potentially partially reduce the size of the bot problem simply by becoming aware of and active in eliminating fraud in their buys.

Conversely, bots and bot fraud traffic patterns evolve and evade in response to scrutiny. Bot traffickers use every available tactic to hide bots among real users, making bot or human determination impossible without the use of bot detection technology.

Other market sectors, such as financial services and retail, have learned through failure and significant losses that periodic audits and certifications are no substitute for continuous and advanced security measures.

To both deter bot traffickers and defend against disguised bots, advertisers must deploy a dual-monitoring strategy: Monitor conspicuously to deter bot traffickers, and also monitor covertly to detect disguised bot traffic.

RECOMMENDATION

To successfully combat bot fraud, advertisers must maintain a public-facing anti-fraud stance and a highly confidential, continuous monitoring program.

Monitor Conspicuously

- Publicly announce your anti-fraud policy to all external partners to temporarily deter certain criminals and fraudsters.
- Announce the intent to conduct audits of all supply-chain partners.

Also Monitor Covertly

- Use bot detection to reveal incentivized human traffic, bot traffic, sourced traffic, and adware in media buys.
- Respond to detected fraud by pausing and troubleshooting campaigns, discussing fraud levels with suppliers, and prioritizing traffic from suppliers that actively suppress and eliminate fraud.

SITES THAT ONLY A BOT COULD LOVE

BOT-TRAFFIC SITES SHOW LITTLE OR NO ORIGINALITY

White Ops visited the study's worst bot-traffic sites. Of the worst 50 sites observed, 30 percent displayed unique content, 22 percent displayed content duplicated within the site, and 51 percent displayed common content (not identical, but highly similar content across pages).

Bot sites copy and paste content.

Of the 50 worst bot-traffic sites:

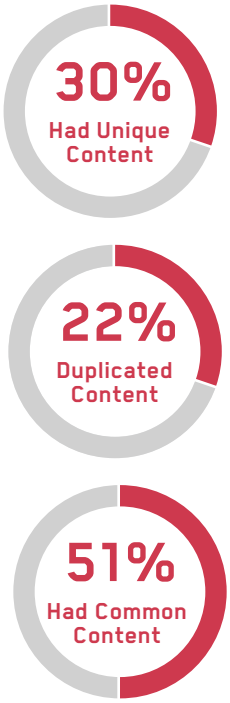

Bot-traffic sites host six or more of the following types of ads:

Video Autoplay

Audio Autoplay

Pop-up/Pop-under

BOT-TRAFFIC SITES DISPLAY MORE ADS

On average, bot-traffic sites hosted six ads per page as opposed to two on comparable top-ranked Alexa sites. Many contained one or more video autoplay, audio autoplay, pop-up, and pop-under elements, which were categorically absent from Alexa's top 50 sites that serve ads.

BOT SUPPLIERS ARE WATCHING YOU

Bot traffickers take careful measurements of traffic to maximize monetezation of their botnets and bot sites. White Ops compared the prevalence of third-party trackers, such as tags, pixels, and beacons, on the worst bot-traffic sites to how often those trackers appeared on Alexa's top 50 sites that serve ads.

The worst bot traffic sites use four times as many trackers

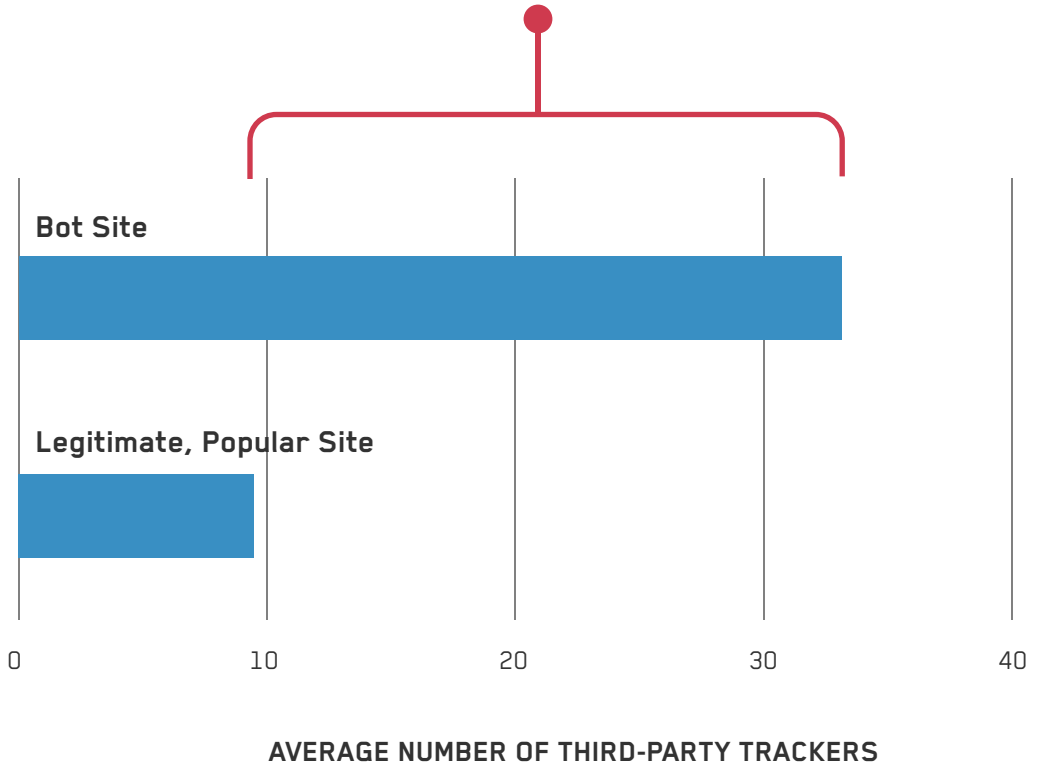

Figure 13: Bot Sites Measure Traffic to Increase Monetization

Tracker measurements were provided by Ghostery.

OLD-SCHOOL BROWSERS HAVE MORE BOTS

Bot percentages soar in older browsers. Development efforts to support older browsers may no longer be cost-effective.

To attack using any new browser, bot traffickers must re-deploy, recompile, and test for each new release. The industry can gain a time advantage in bot defense by optimizing media for the newest browsers while reducing impressions from obsolete browsers.

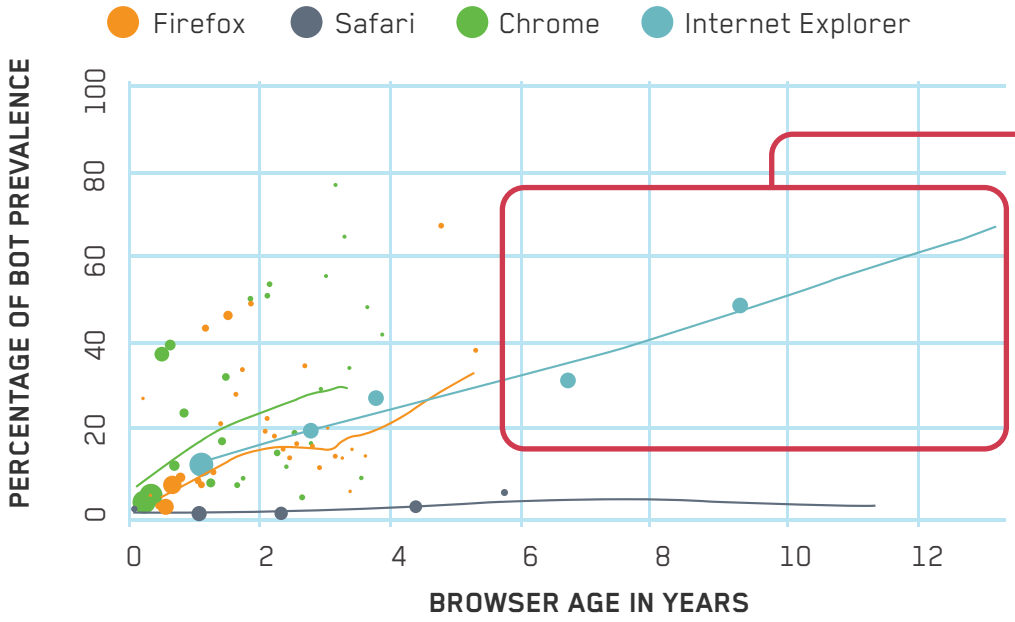

FINDING

Over half of the impressions (58%) coming from IE6 browsers and almost half of those (46%) from IE7 browsers are bots.

Figure 14: Bots are Still Using IE6 and IE7

RECOMMENDATION

Bot browsers do not typically auto-upgrade the way popular browsers auto-upgrade for human users. Targeting media buys to newer browsers can increase bot downtime and cost while enabling the industry to get ahead in the bot-defense timeline.

Support Newer Browsers.

- Prefer newer releases to older ones.
- Know which browser versions are more bot than human.
- Save on development costs by reducing support for obsolete browsers that do not deliver significant human audiences.

WHEN PUBLISHERS ARE VICTIMS TOO: AD INJECTION

AD INJECTION ATTACKS USERS, HARMS PUBLISHERS, AND DEFRAUDS ADVERTISERS

Ad injection attackers use the same mechanisms that cybercriminals use to rob online banks. Payment and impression data for injected ad inventory flow to third parties which have no affiliation with the sites on which the ads are displayed. Botnet controllers who inject ads are able to generate enormous revenues using content and brands they did not build or maintain.

Ad Injection

A man-in-the-browser attack on publishers, advertisers, and users in which ads are forced onto a website, often displacing the initial web page content or overlaying on top of existing content or ads

We did not set out to detect ad injection in this study. However, using White Ops' true domain detection technology, **we found significant evidence of ads running on sites which are well known as user-funded or subscription-based sites that do not permit ads.** These injected ads were unsanctioned by the publishers. The ads were displayed using malware illicitly installed on residential computers.

Victims of malware-driven ad injection inadvertently expose private information. Malware-driven ad injection software on the victim's computer allows potentially malicious, unknown actors to gain access to personally identifiable information (PII), including browsing history, interests, and financial information.

True Domain

Technology that identifies the actual domain on which an ad displays rather than the domain reported by the ad server which can be falsified

CASE STUDY

Over 500,000 Ads Were Injected Daily at One Publisher

500,000 Per Day

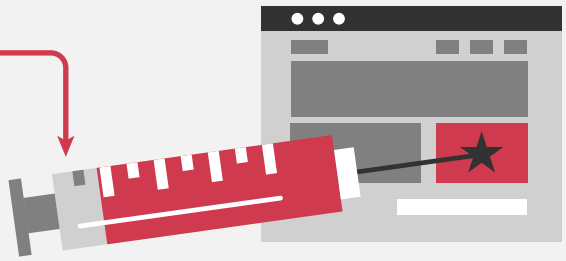

Post-study analysis showed that a single publisher was the victim of a minimum of 500,000 injected ads per day through the duration of the study.

AD INJECTION DEGRADES THE QUALITY OF THE INTERNET

The degradation of the user's browsing experience due to ad injection includes:

- Websites with injected ads load more slowly.
- Too many ads on the page can overwhelm the user.
- Injected ads are often highly intrusive and distracting and may break the intended functionality of the displayed website.

AD INJECTION DESTROYS LEGITIMATE AD INVENTORY

Advertisers and publishers do not choose to inject ads into a site. The owner of the ad injection malware accepts payment for unsanctioned media that could potentially damage both the advertiser's and publisher's reputation and deplete the advertiser's digital ad inventory budget.

Ad injection also devalues all authentic advertising running on the site.

Advertisers and publishers can help prevent ad injection by limiting the use of sourced traffic, continuously monitoring sourced traffic, and requiring suppliers (including DSPs) to demonstrate that their traffic does not include injected ads.

ELIMINATING BOT FRAUD: A CALL TO ACTION

BOTNET OPERATORS ARE ALREADY DEFEATING SOME COUNTERMEASURES

Several tactics that agencies and ad-tech platforms may currently regard as effective in preventing bots were mostly ineffective:

- Technical measures of viewability do not ensure humanity. Since so many bot operators have upgraded their bots to fake viewability, viewable impressions actually skewed slightly higher in bot incidence than non-viewable impressions.
- Blacklists require near real-time updating and often block significant volumes of real human audience as well as bots. Fraudsters adapt quickly.
- Optimizing campaigns with even the most sophisticated engagement metrics and attribution models did not eliminate bot traffic, because bots clone real people (getting bots credited for purchases made by real people), and bots fake engagement.
- Due to the pervasiveness of traffic sourcing, buying strictly from premium publishers did not eliminate bot traffic.

Cui Bono: Who Benefits?

Bot impressions distort the entire market by making it look as though there are more people viewing ads than there really are. The illusion of an unlimited, diverse supply of ad inventory drives the price of real human impressions down. It puts honest players at a huge competitive disadvantage, pressuring them to source traffic, too.

The motivations and temptations of various parts of the ecosystem differ significantly:

- Botnet operators extract their payments through cash-out points, the final stop in the fraud supply chain.
- Aggregators and middlemen gain reach, ensuring they never lack inventory to sell, and a diversity of bot profiles that match any conceivable audience segment.
- Publishers inflate their apparent audience size and pocket the difference between their traffic acquisition cost and the revenue received from advertisers.

FEAR IS THE UNSPOKEN OBSTACLE

The ecosystem described in this report is complex. There are winners and losers in advertising fraud today, and the scoreboard clearly is tilted in the wrong direction.

On one side, the winners are raking in billions of dollars, much of which funds cybercriminal activities by bot suppliers who have no incentive to change their behavior. Far behind in the digital advertising game are advertisers who want to offer great products to the right customers, agencies who want their media plans to reach the appropriate targets, publishers who want to support the content on their sites through pertinent advertising, and the advertising technology community who want to provide an innovative infrastructure and marketplace for online advertising.

Also losing big are consumers, who are the real reason the digital ad industry exists and who have been turned into unwitting accomplices in vast networks of botnets.

The common reaction to the entire issue of fraud is fear. Fear leads to avoidance, which is just what the bad guys want. **No one wants to look unaware, unscrupulous, or negligent enough to "allow" this kind of activity to take place.**

The truth is: fraud is everywhere. No one is immune. Only by emancipating your people and partners from that fear can we get the cooperation needed to address this issue effectively.

Bot fraud is a new type of attack. Advertisers, agencies, and publishers must learn and use new concepts to understand the bot fraud threat that has emerged over the past few years. Advertisers, and all industry participants, can and must take action. Some actions can be taken unilaterally; others must be done in partnership with a fraud detection partner. The following pages provide an action plan for the stakeholders in the industry to combat fraud in digital advertising.

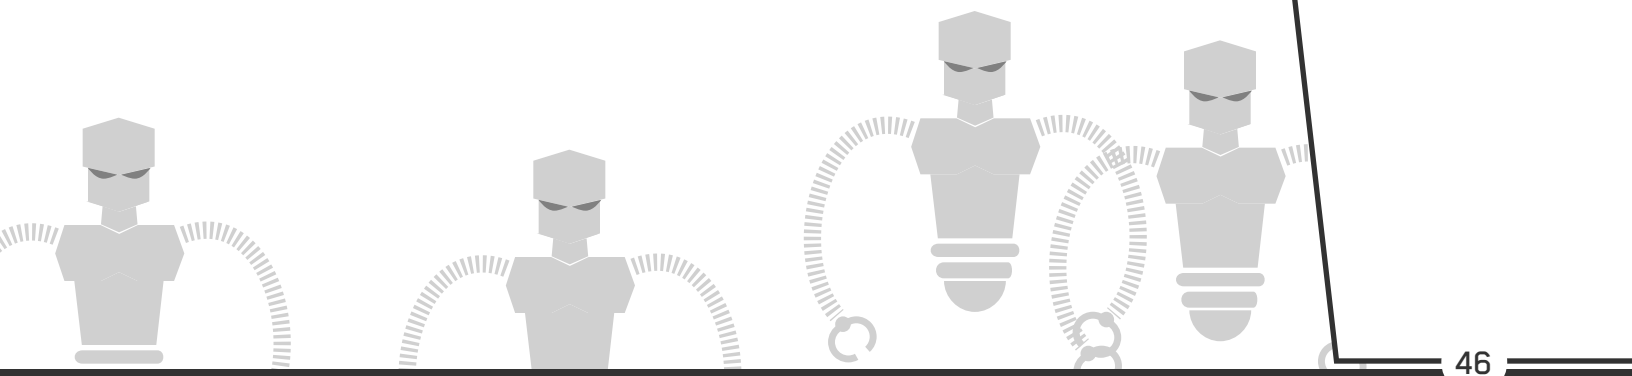

Action Plan for All Stakeholders

Create allies, not adversaries, in the fight against bot fraud

Bot fraud affects many suppliers in the digital advertising supply chain before it reaches the advertiser. Advertisers, agencies, and suppliers must all work cooperatively to reduce and eliminate bot fraud in the supply chain. Our call to action is for the key industry players to work both collaboratively and individually to substantially reduce bot fraud.

Manage the emotions of ad fraud discussions

Do not assume that bot fraud in your campaigns indicates an agency or publisher is deficient or bad. Remember, it's likely that your media seller is a victim of the botnet operators, not the cause.

Authorize and approve third-party traffic validation technology

This study was not deployed across all participants' placements, partly due to agency and publisher policies. Some agencies and publishers did not permit the monitoring software in certain placements (see *Appendix B: Constraints and Limitations,* page 55).

To effectively combat bots in their media buys, advertisers must be able to deploy monitoring tools. Publishers and agencies must enable the deployment of these monitoring tools. Set policy and procedures to enable advertisers to deploy bot detection and domain detection software to their ad buys.

Support the Trustworthy Accountability Group

The IAB, 4A's, and the ANA announced in early November the creation of the Trustworthy Accountability Group (TAG), a joint marketing-media industry program designed to eradicate digital advertising fraud, malware, ad-supported piracy, and other deficiencies in the digital communications supply chain. All vendors should comply with TAG's quality assurance guidelines.

While Taking Actions Against Bot Traffic, Communicate About Bots Effectively:

Within your organization, use language that accurately communicates the bot fraud problem.

Add bot-fraud discussion time to all media buy conversations internally and externally.

Adopt and use terms that correctly identify threats and real adversaries while preserving allies and building an alliance against fraud.

Action Plan for Buyers

Be aware and involved

Advertisers must be aware of digital advertising fraud and take an active and vocal position in addressing the problem. Fraud hurts everyone in the digital communications supply chain, especially advertisers. Advertisers must therefore play an active role in effecting positive change.

Request transparency for sourced traffic

Traffic sourcing correlates strongly to high bot percentages. It's recommended that buyers request transparency from publishers around traffic sourcing and build language in RFPs and IOs that requires publishers to identify all third-party sources of traffic. Furthermore, buyers should have the option of rejecting sourced traffic and running their advertising only on a publisher's organic site traffic.

Include language on non-human traffic in terms and conditions

Consider adding specific language to your terms and conditions to address the issues discussed in this study. An illustration of one approach to the definition of fraudulent traffic and the safeguards that might be negotiated between advertisers and media companies is provided in the appendix (developed by Reed Smith, ANA's outside legal counsel). You should consult with your own counsel to develop specific provisions that best serve your company's individual interests (see *Appendix D: Illustrative Terms and Conditions*, page 57).

Use third-party monitoring

Monitor all traffic with a consistent tool. Comparability is essential. Selective monitoring, such as once a month, once a quarter, or only on certain channels, encourages evasive maneuvers by bot suppliers. Third-party monitoring can validate or disprove assumptions about the quality of a publisher or ad tech company's traffic. We recommend relentless monitoring to get the best value out of your ad investment.

Use monitoring and bot detection to reveal the bots in retargeting campaigns and audience metrics. This will prevent the purchase of additional media targeted at those bots and will improve campaign metrics.

Action Plan for Buyers (Continued)

Apply day-parting when you can

Bot fraud represents a higher proportion of traffic between midnight and 7 a.m. Buyers can reduce bots by concentrating advertising during audience waking hours.

Update blacklists frequently and narrowly

Be careful how you block. For blacklists to be effective, they need to be updated at least daily, must be very specific (micro-blacklisting), and must accompany other defenses.

Control for ad injection

Ad injection (the unauthorized placing of ads on sites where they do not belong) is a tactic that causes programmatic buys to contain higher levels of fraud. Discuss with your DSP or tech platform how to control ad injection.

Consider reducing buys for older browsers

There are more bots claiming to be IE6 (2001 original release date) or IE7 (2007 original release date) than there are real humans still using those browsers. Consider reducing older browser impressions in buys.

Announce your anti-fraud policy to all external partners

In combination with covert, continuous monitoring practices, the watchdog effect will change behavior, reduce fraud, and encourage others to join the fight.

Budget for security

Across many industries, the typical cost of security amounts to an overhead of 1 to 3 percent. In the credit card ecosystem, that security spending has lowered the losses due to fraud to just $0.08 cents per hundred dollars. Lowering bot fraud in advertising to those levels could potentially return many multiples of the security spending needed to achieve it.

Action Plan for Publishers

Continuously monitor sourced traffic

Always monitor sourced traffic. Know your sources and maintain transparency about traffic sourcing. Eliminate sources of traffic that are shown to have high bot percentages. Monitor all vendors, all the time.

Protect yourself from content theft and ad injection

Use a service such as domain detection or bot detection to monitor for content-scraping (presenting another site's content in a separate website and monetizing the scraped content with ads) and evidence of ad injection. A bot detection service can measure actual numbers of bots in high-bot traffic, allowing payment for the human audience while eliminating bots from the billing process.

Consider allowing third-party traffic assessment tools

Publishers can enable advertisers to improve the granularity of their traffic performance by authorizing third-party monitoring (for characteristics such as viewability, engagement, and bot detection) and third-party tracker measurement.

Who is your **real** *audience:*

BOB SMITH or BOT SMITH ?

Thank you and congratulations on becoming part of the fraud elimination movement.

We encourage you to provide "protected transparency" to all your digital media partners in order to encourage collaborative honesty on behalf of your brands and the industry.

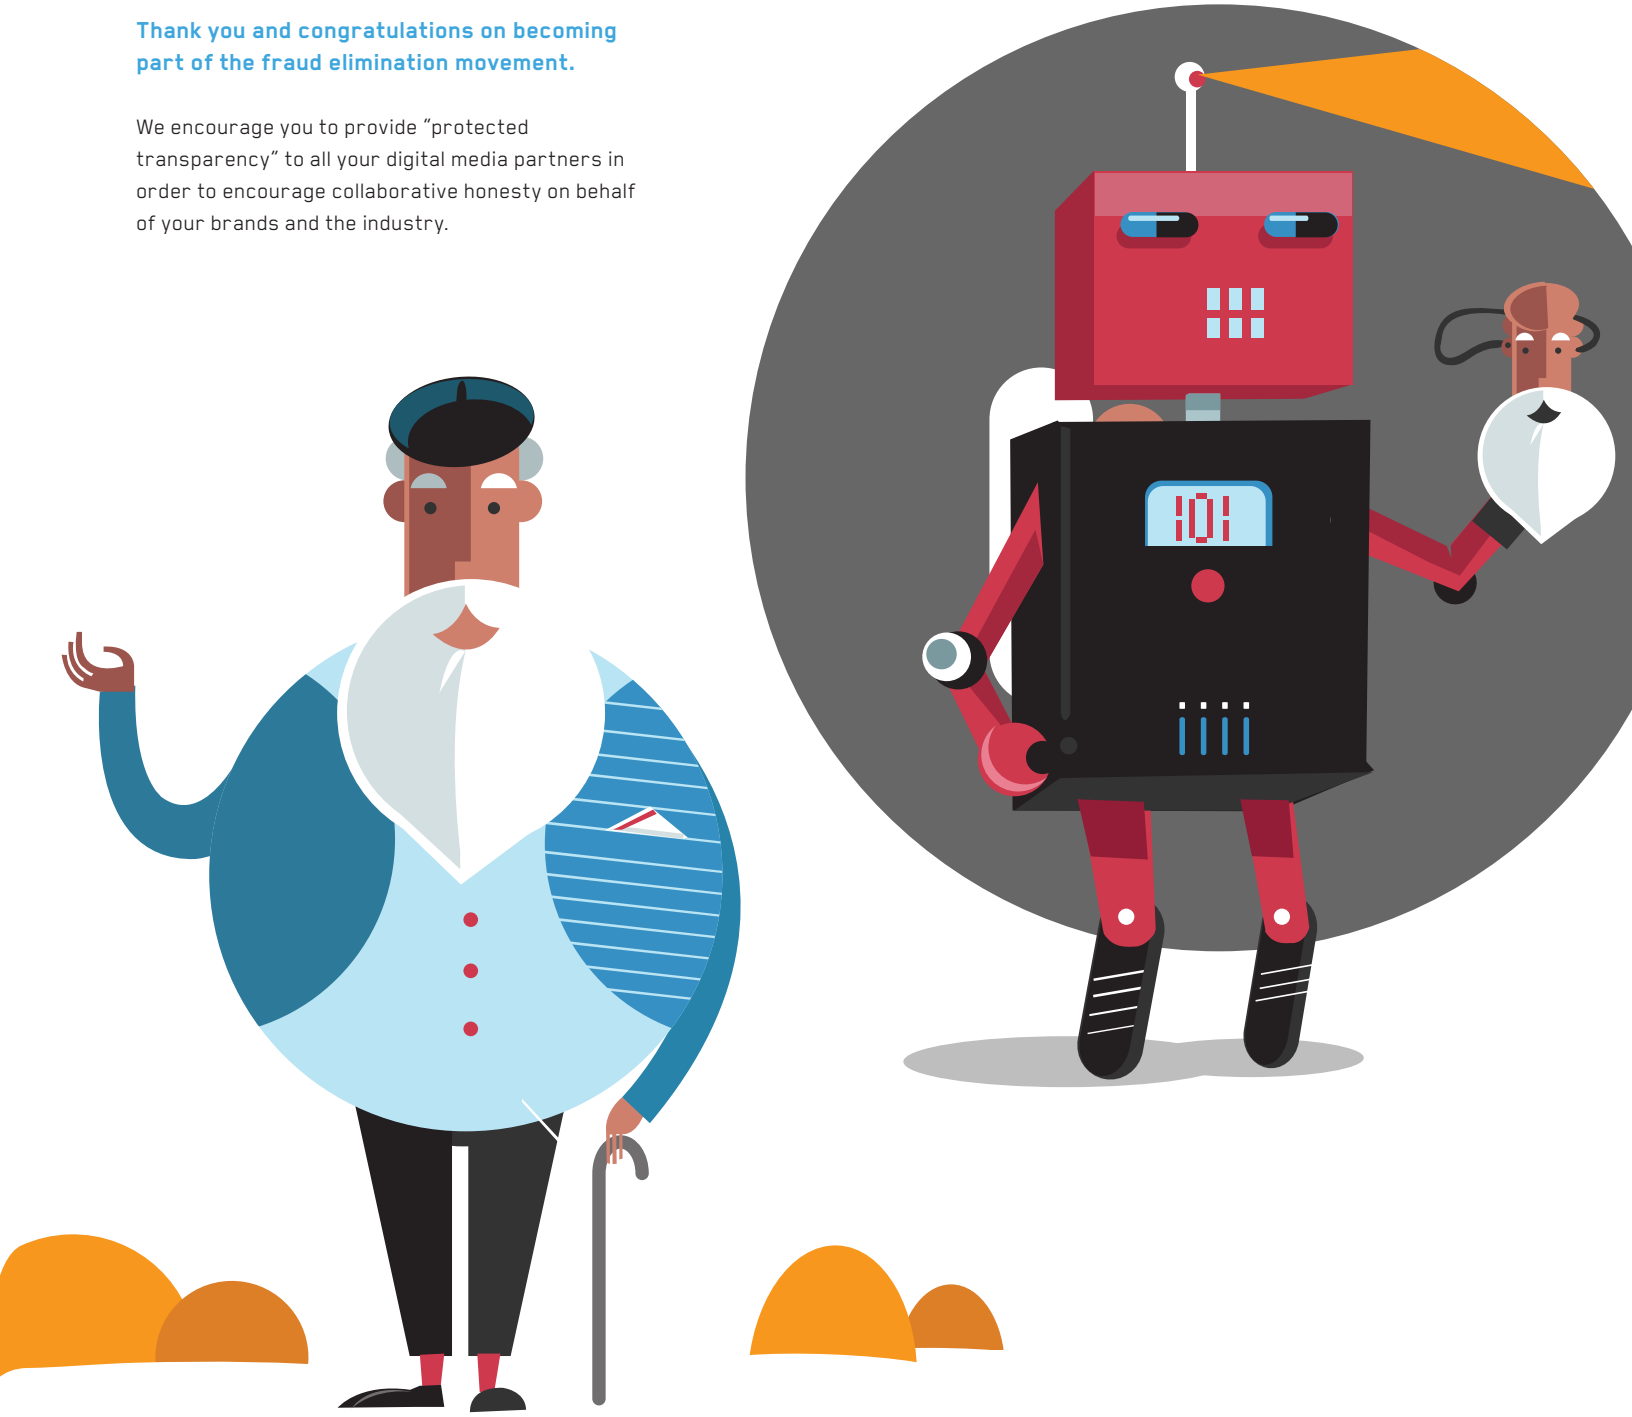

Appendix A

Glossary of Terms

Ad
An online advertisement of any sort

Ad Fraud
The inclusion in reports, bills, or other analytics of anything other than natural persons consuming **ads** in the normal course of using any device

Ad Injection
The visible or hidden insertion of **ads** into an app, **web page**, or other online resource without the consent of the **publisher** or operator of that resource

Advertiser
A company, brand, or individual who pays a third party to display or act as agent for the display of **ads**

Adware
Software, often automatically installed on user devices, that displays visible or hidden **ads** to **users** to boost **ad** consumption

Autoplay
The playing of a sound, video, or any other type of media, generally as part of an **ad**, without any **user** interaction, often upon the **user's** loading of a **web page** or other resource

Bot(s)
(a/k/a Non-Human Traffic or NHT) Automated entities capable of consuming any digital content, including text, video, images, audio, and other data. These agents may intentionally or unintentionally view ads, watch videos, listen to radio spots, fake viewability, and click on ads

Bot Percentage
The percentage of a given portion of online advertising **traffic** consumed by **bots**

Bot Detection
The detection and differentiation of **bot traffic** and **bot impressions** from **human traffic** and **human impressions**

Bot Fraud
Ad fraud specifically perpetrated by **bots**

Bot Impression
An **impression** consumed by a **bot**

Bot Traffic
Automated website or other online **traffic** and/or **ad** consumption driven by or resulting from **bots**

Botnet
A group of infected computers that generate automated web events. The infrastructure used to create many types of **bots**

Botprints
A unique combination of directly observed properties in a given **impression**, page view, or other online event which collectively identifies that event as **bot**-driven by a specific type of **bot**

Broker
Third-party arbitragers that buy traffic from **suppliers** and sell to **publishers**; often media agencies, retargeting platforms, or traffic extension platforms

Campaign
A series of **ads** on behalf of an advertiser that share a single idea and theme, and which may be made up of different types of **ads**, and which may be run on multiple **publishers**, **sites,** or other channels and in multiple formats

Campaign Monitoring
Monitoring the various types of **ads** and their formats, and the publishers, **sites,** and channels on or in which they are displayed, for the purpose of detecting differing levels of **ad fraud**, allowing for the optimization of spending to reduce **ad fraud**

Cash-Out Site
A website, app, or other resource that is capable of delivering **ads,** and is operated by perpetrators of **ad fraud** for the purpose of exfiltrating money from the online advertising ecosystem

Appendix A

Glossary of Terms (Continued)

Click Farm
A type of **ad fraud**, in which a large group of human workers (in one or multiple locations) is minimally paid or otherwise incentivized to view and/or click on **ads** on behalf of a third party that economically benefits from those human workers' illegitimate consumption of those **ads**

Decision
The deterministic, evidence-based identification of a particular **impression**, **page view,** or other type of online event, either legitimate or the result of **ad fraud**

Domain
A unique name that identifies and can be used to access an Internet resource such as a website

Domain Blacklisting
Using lists of known bad **domains** to prevent the serving of **ads** to those **domains**

Domain Detection
Determining the **domain** on which an **ad** was actually displayed, as opposed to the domain which an ad server may report

DSP (Demand-Side Platform)
A platform that allows advertisers or their agencies to manage multiple **exchange** accounts and bid across those accounts

Engagement
A metric (often defined with great specificity) that provides a qualitative evaluation of a **user's** interaction with a given **ad** or web page

Exchange
A technology platform that facilitates the buying and selling of **ads** and related data from multiple sources such as **publishers** and networks of **publishers**

HREF Domain
The full domain path representing the location where a particular **impression**, **page view,** or other online event occurred; often forms part of **ad** serving reports

Human Impression
An **impression** legitimately served to a real human not intentionally or unintentionally engaged in any form of **ad fraud**

Human Traffic
Legitimate website or other online traffic and/or **ad** consumption driven by real humans

Impression
A particular instance of the delivery of a particular online **ad**. The basic economic unit of online advertising, generally as recorded by ad servers for the purposes of billing **advertisers** or their agencies

Incentivized Human Impression
An **impression** served to a human who is paid or otherwise incentivized

IP (IP Address)
A unique numerical address corresponding to a particular device or set of devices connected to the Internet

IP Geolocation
Determining the approximate physical location of a device connected to the Internet at a given point in time by using information associated with or deduced from that device's **IP address**

IP Blacklisting
Using lists of known bad **IPs** to prevent the serving of **ads** to those **IPs**

Long Tail
Websites with relatively low traffic that may offer value to **advertisers** due to their appeal to specific or niche audiences of **users**

Make-Good
Credit given to an **advertiser** or their agency (or the use of that credit) to compensate for an error in the composition, placement, or delivery of an **ad**

Appendix A

Glossary of Terms (Continued)

Man-in-the-Browser Attack
An Internet attack that infects a **user's** online interactions by taking advantage of vulnerabilities in browser or app security to modify **ads**, web pages, or transaction content or to insert additional **ads**, content, or transactions, without the knowledge or consent of the **user** or the resource(s) with which the **user** intended to interact

Micro-Blacklist
A blacklist that is updated and expires frequently, to enhance its effectiveness against advanced and adaptive threats

Page View
A single request to load a single page of a website

Phantom Layer
Websites operated specifically for the purpose of laundering **ad fraud** by obscuring the source of inventory and impressions entering the online advertising ecosystem

Pop-Under
Windows that appear or open under the **user's** current browser window so that they become visible when that window is closed

Pop-up
Windows that appear or open above or on top of the **user's** current browser window

Publisher
The operator of a website or network of websites, and the producer or curator of content for those sites. A seller of online advertising space and impressions, and often a buyer of third-party traffic

Reach
The total number of different **users** exposed, at least once, to an **ad** or campaign during a given period of time

Retargeting (Behavioral Retargeting)
The process of delivering **ads** to particular **users** based on their previous online activity

RON (Run-of-Network)
An **ad** or campaign displayed on a large collection of websites without the ability to choose target-specific sites, placements, or **domains**

Site or Website
A set of related web pages, often served from a single **domain**

SSP (Supply-Side Platform)
A technology platform that enables publishers to manage their **ad** inventory and maximize revenue from online advertising, usually by interfacing with **ad** exchanges, and making their **ad** placement inventory available in an automated fashion to a wide number of potential purchasers

Supplier
A seller of traffic to **publishers** and sites

Traffic
Visits to a particular site, page, or other online resource; impressions related to a particular **ad**

Traffic Sourcing or Sourced Traffic
Any method by which publishers acquire more visitors through third parties

True Domain
The domain on which an ad actually ran, as determined by domain detection

User
A person who uses a computer or other device or network service. In the context of online advertising, a visitor to a **publisher's** site, and a consumer of an advertiser's **ads**

Appendix B

Constraints and Limitations

Complexity of Study

Study participants joined the initiative for different periods of time with varying platform configurations, target audiences, industry verticals, and ad agencies. Because of these differences, not all data from the study can be compared directly between participants.

Web Framework Limitations

The study software could only be deployed to systems that were JavaScript-enabled.

Public Study Awareness

Because many participants experienced administrative and technical deployment delays during the public study phase, and because bot numbers may have been artificially decreased during the public study phase, bot numbers detected during the covert study phase may be more representative of the numbers occurring in normal ad campaigns outside of the initiative framework. Because the study was announced and widely known, it is assumed that bot numbers for the month of the public study were artificially lower than numbers we might otherwise observe.

Seasonal Time Frame

White Ops expects bot numbers to be at their lowest in the late summer months and at their highest when demand for advertising is highest near the end of the calendar year. This seasonal snapshot from the months of August and September cannot predict the number of bots in a typical month or during peak months when advertising volume is higher.

Coordinating With Agencies

Some study participants were not able to deploy the study software uniformly throughout their campaigns due to administrative elements including legal agreements, site policies, and organizational complexity. In some cases, the study participants were not aware that study software had not been deployed through their ad agencies. When these issues became apparent during the monitoring and data collection phases, White Ops worked with the study participants and their ad agencies to attempt to correct the problem.

Appendix C

External Study Contributors

White Ops worked with several study contributors to provide additional insight into the study results.

Chartbeat

Chartbeat, a betaworks company that provides real-time analytics to websites and blogs, provided data for comparison of bot and human engagement and viewability metrics. Chartbeat matched 120 million impressions with White Ops data on 87 publishers.

Ghostery

Ghostery® is a global technology company that provides solutions for online transparency and control to individuals and businesses. Ghostery provided insight into trackers running on the study's top 10,000 domains by traffic volume, the study's worst bot sites, and top Alexa sites.

Grapeshot

Grapeshot is a software technology company using advanced Information Retrieval techniques pioneered at Cambridge University. For the study, Grapeshot provided domain characteristics data for 13,000 domains. White Ops used this data to identify bot trends by domain category.

Appendix D

Illustrative Terms and Conditions

Consider adding specific language to your terms and conditions to address the issues discussed in this study. An illustration of one approach to the definition of fraudulent traffic and the safeguards that might be negotiated between advertisers and media companies appears below (developed by Reed Smith, ANA's outside legal counsel). You should consult with your own counsel to develop specific provisions that best serve your company's individual interests.

T & C

Fraudulent Traffic

(a) "Fraudulent Traffic" means the inclusion in reports, bills or other information and materials associated with this Agreement, of data that counts or uses in calculations, anything other than natural persons viewing actually displayed Ads in the normal course of using any device, including, without limitation, browsing through online, mobile or any other technology or platform. For the avoidance of ambiguity, Fraudulent Traffic includes, without limitation, the inclusion or counting of views: (i) by a natural person who has been engaged for the purpose of viewing such Ads, whether exclusively or in conjunction with any other activities of that person; (ii) by non-human visitors; (iii) combinations of displays directed or redirected by any combination of (i) and/or (ii); and (iv) that are not actually visible to the human eye, discernible to human senses or perceived by a human being.

(b) Media Company will establish, implement and use all commercially reasonable technology and methodologies to: (i) prevent Fraudulent Traffic; (ii) detect Fraudulent Traffic should it occur; and (iii) promptly take steps to prevent continuation and/or recurrence of occurrences thereof. Media Company will ensure, by agreement, instruction or any other legally enforceable means, that all third parties to which Ads are delivered, displayed or made available have adopted and implemented technology and methodologies (and agreed in writing thereto) to ensure Media Company is in compliance with the foregoing obligations. Media Company agrees that Advertiser shall have no obligation hereunder, for compensation, liability or otherwise in respect of Fraudulent Traffic and shall not be billed or required to pay for Fraudulent Traffic. To the extent any payment attributable to Fraudulent Traffic is or may be paid by Advertiser, Media Company shall, within five (5) days, reimburse and refund such payment to Advertiser, together with reasonably adequate documentation to substantiate the accuracy of any such reimbursement or refund. Unless otherwise included in another audit provision hereunder, Advertiser or its designated auditors shall be entitled to audit the books and records of Media Company for the purpose of determining compliance with the provisions of this Agreement.