The Fraud is Coming from Inside the Ads

DEMAND-SIDE AD FRAUD AND WHAT IT COSTS LOCAL MEDIA

published for the LMA Revenue Summit, April 2019

It's Saturday morning, January 23, 2019. The gossip at a New England coffee shop turns to a downtown zoning dispute, and a woman pulls out her iPhone to check a local TV news site for the latest developments. Her location, the subject of the article, and the length of her visit are tracked by a data aggregator in Turkey.

It's mid-afternoon on Monday, February 5. A sheriff's deputy in an Atlantic coast town smiles as she reads her quote in a story about a breakthrough on a cold case. An Iraqi traffic counter logs her visit.

It's early evening on Sunday, March 17. A man named Greg is sitting at his computer in a Midwestern college town, furiously typing a comment about the NCAA selection committee's lack of intelligence. Ads load on his favorite local sports site – and so does a tracking pixel from a Russian server farm.

These incidents aren't isolated, nor are they uncommon. Research conducted across more than 1,000 local media sites monitored more than 16 million cases where scripts from foreign servers were served on US local news sites in the first quarter of 2019. And, contrary to popular thought, it did not occur because of so-called "fake news" sites designed to spoof, hoodwink or simply distract news consumers.

It happened in the ads served on those pages.

This was not the fault of the revenue and sales departments – local or corporate – that oversaw filling the ad slots that carried the trackers. Nor was it a technical or administrative breakdown by the publisher that inadvertently exposed its users to foreign servers. Nor was it the fault of any of the hundreds – often thousands – of ad networks providing programmatic ads to meet audience demand. Those are the facts, although all players in the events leading up to an incidence of ad fraud will point fingers at each other.

All of them are at fault, because digital ad fraud is a systemic problem in the business of online media. Each player in the ecosystem bears some responsibility for news consumers that are now being tracked by far-flung data collectors, because that is exactly how the system is supposed to work. It's also its greatest weakness.

Bad actors in ad fraud target only one thing: Weakness. In a system where the buck stops nowhere, there is ample opportunity for fraud, money laundering, ransomware, watering hole exploits, browser hijacking – all of it.

It's how the system works.

Juniper Research made a splash in 2017 with its conclusion that digital ad fraud would be a $19 billion issue and would grow to more than $40 billion by 2022. The buy-side of the industry – major advertisers, brands and agencies – have loudly and rightfully demanded that digital content publishers provide third-party verification that their ads were seen by humans, not bots, that the ads were sufficiently viewable and were not blocked by an army of browser plug-ins.

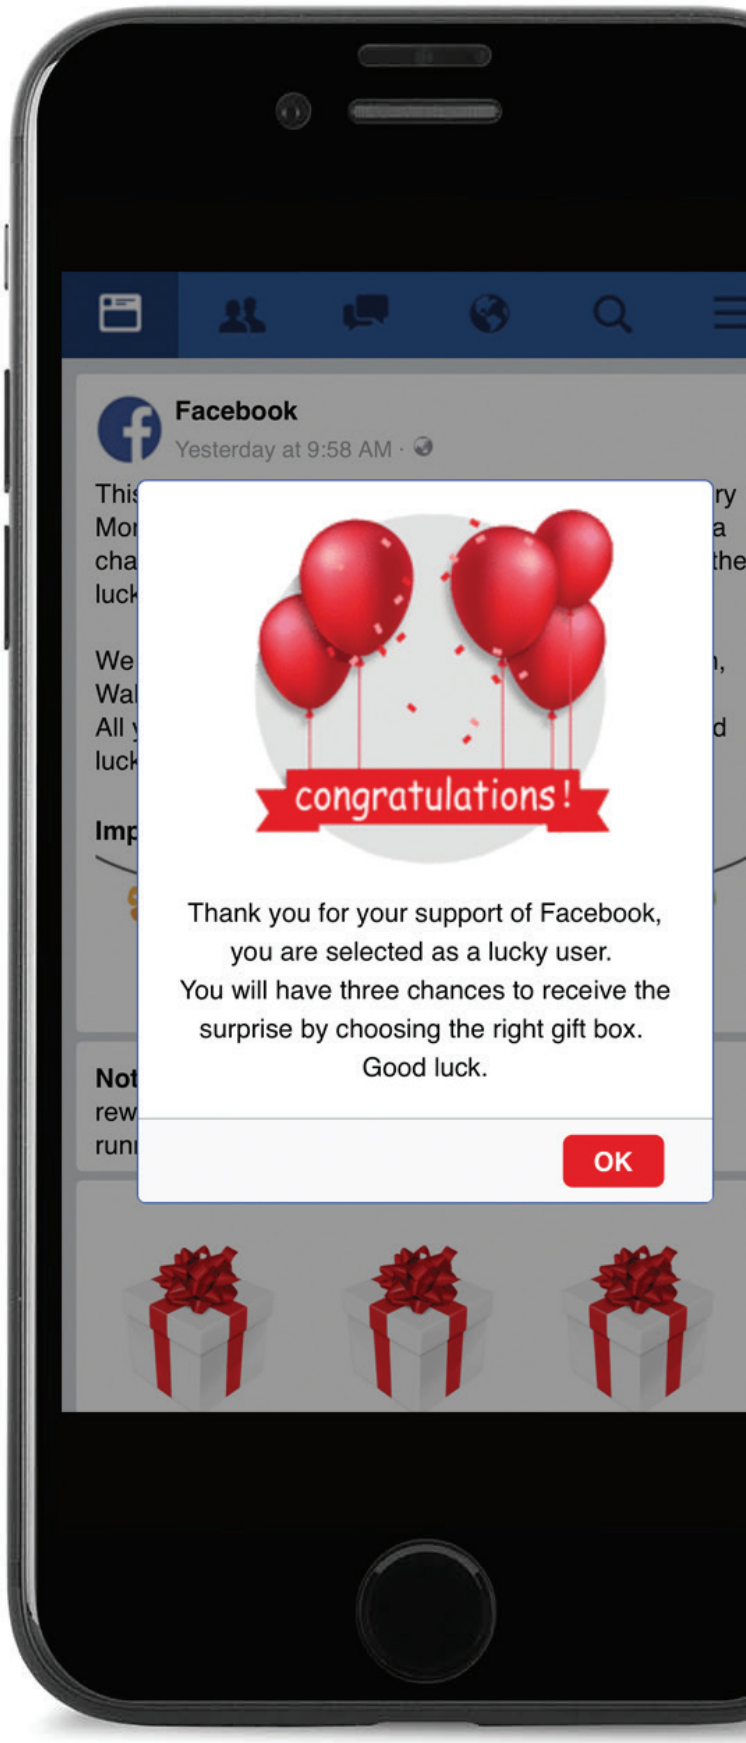

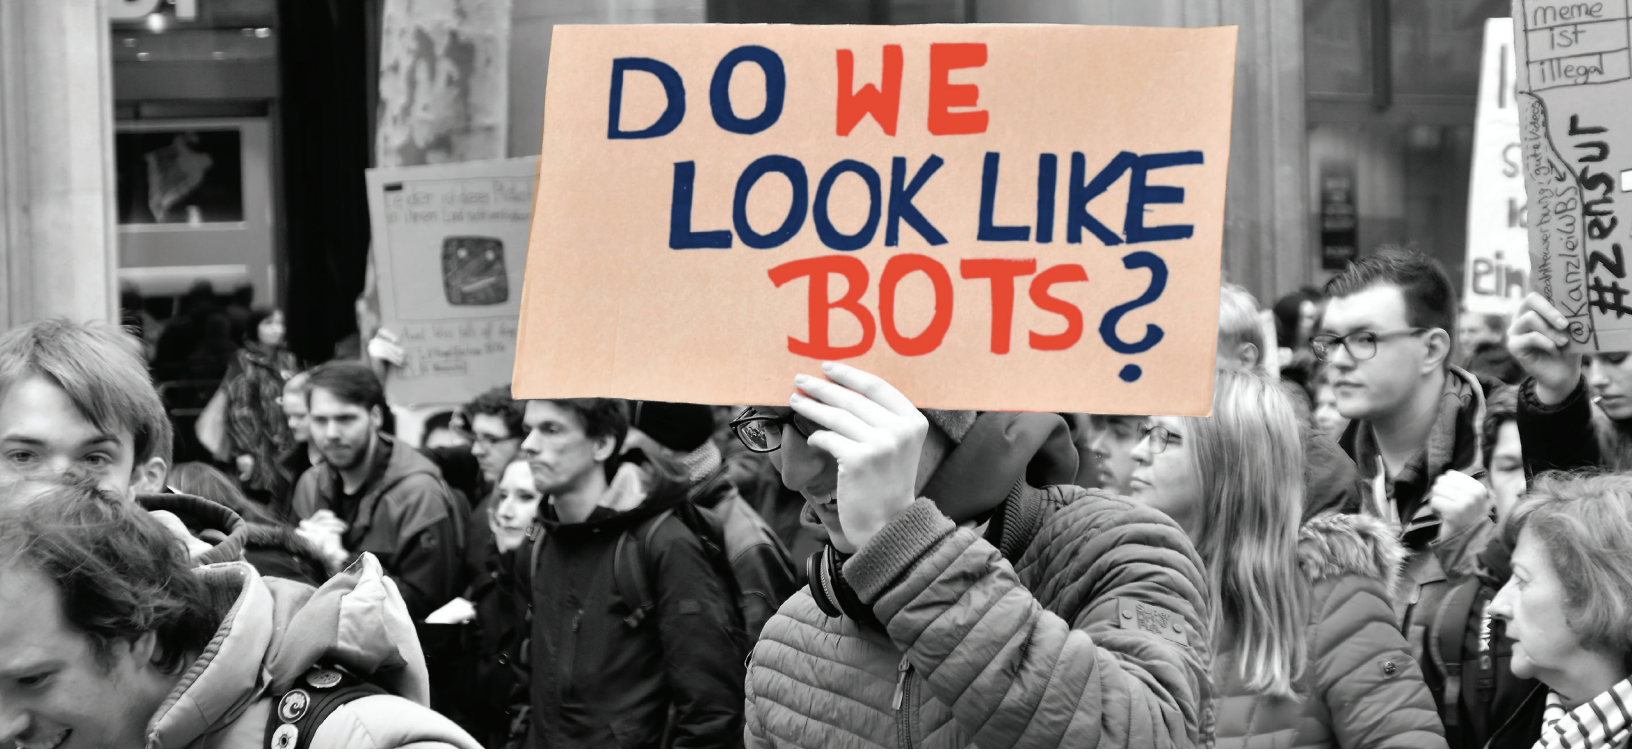

That's where the ad fraud conversation usually ends – the supply-side problem of invalid traffic. But that's only one side of the fraud equation. Criminals use the inviting advertising ecosystem to defraud users, earn income and launder money by injecting deceptive creative, browser-hijacking pop-up ads and malware. While publishers certainly need to invest in certifying their impressions, they also need to protect their readers, their ad revenue and their brand's reputation from these demand-side threats.

Publishers – in this case, local news publishers, including newspapers, broadcasters and digital natives – have been beset with these previously unbudgeted costs, further cutting their margins on what already is a challenging business. Because while it's well publicized that the major platforms – Google, Facebook Amazon and Oath – bring in the lion's share of all digital advertising revenue, those left to scrap for what's left must do so in ways that increasingly have the potential to expose users to risk.

Again, it's how the system works.

To understand the depth of the issue at the local news level – where, according to the recent Knight Commission report "Crisis in America: Renewing Trust in America" is where trust in news is most likely to be built – AdHack.org formed a partnership with the Local Media Association, ad fraud detection company DEVCON, ad operations and yield provider OpsCo to offer free ad fraud detection tools to all local media outlets in the LMA's membership. This allowed for several thousand local publishers – again, local newspapers, TV stations, radio stations and digital natives – to monitor the extent of ad fraud occurring on their sites at no cost. The initiative is called Freedom for Media.

Even without a price, the proposition was not a slam-dunk, with implementation being turned down because of technical development schedules, other priorities and sometimes the opinion that digital ad fraud on local news sites affects less that 1 percent of a site or company's ad impressions. "It's a rounding error," one Chief Digital Officer said at the outset of the study. "I don't see the ROI."

The first 90 days of the study proved how incorrect – and toxic – that opinion is.

In the first 90 days, the AdHack.org study saw 4 billion ad impressions flow through the system of more than 1,000 local media sites. And in those 4 billion impressions, DEVCON's researchers identified 16.5 million exploited ads.

Let's put that another way. If, as the brands and agencies require, each of those ads was seen by a human being, that means 16.5 million local news consumers were affected – or infected – by bad ads the publishers likely had no idea were flowing through their trusted news sites.

Sixteen point five million people. Getting local news.

Exploits come in hundreds of varieties – some benign (they just want to better target users) to criminal (they want to take over the user's desktop or phone). In the first 90 days of testing, the AdHack.org study detected the following exploits.

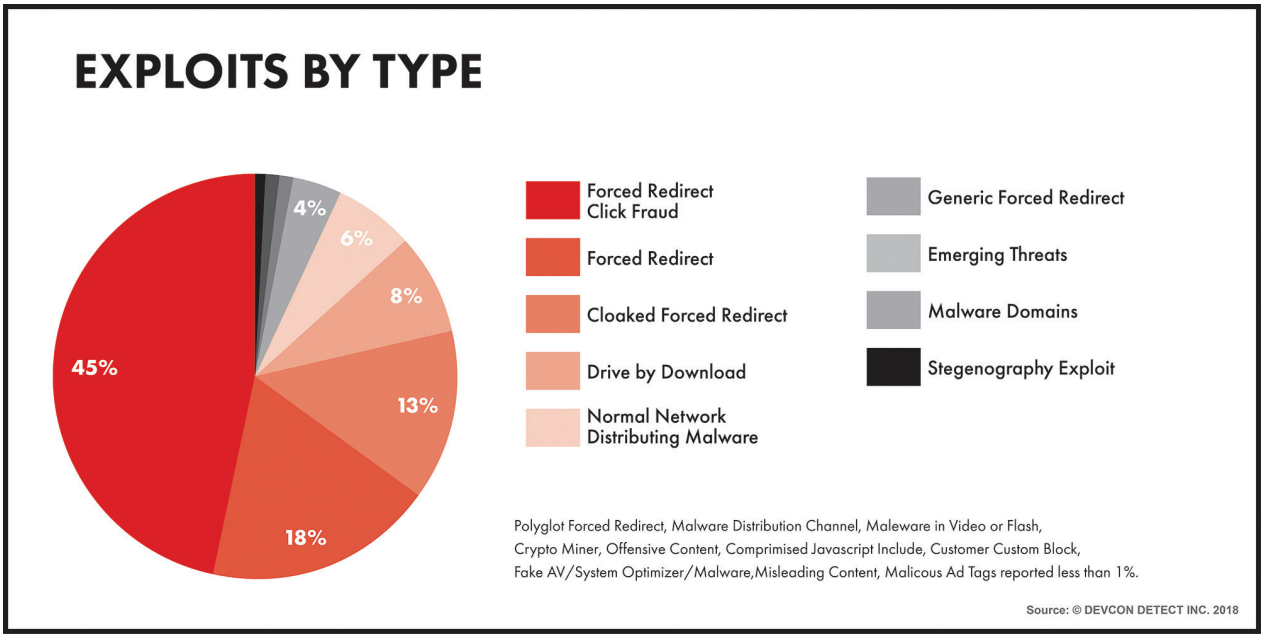

Those exploits accounted for .45 percent of all ad impressions across the test group – less than 1 percent (which may be considered a rounding error by some) but representing 16.5 million total impressions (which is only a rounding error in astronomy). The research also found that 72 percent of those exploits were forced redirects and 22 percent were other types of server exploits.

Forced redirects – in which the user is, by no action of her own, sent to another website (or a series of websites in short succession) – cause a variety of digital ills. Many take users through a cascade of fraudulent websites filled with pirated content and surrounded by an arcade-like abundance of low-CPM ads, with the intent of driving up "views" on the ads so the site owners make easy money. Others send users to false roulette or trivia games, touting that Amazon, Walmart or others are offering gift cards of $1,000 or $10,000 – as long as the user enters a seemingly endless array of personal and behavioral data.

Spoiler: There is no gift card at the end of this journey.

"Cyber threat actors are just like the rest of us: they want to succeed, but their ethics are questionable, to say the least," said Michael F.D. Anaya, Head of Global Cyber Investigations and Government Relations at DEVCON. Anaya formerly was Supervising Special Agent for the FBI. "Success in their world is just like it is in our world: they set a goal and when that goal is accomplished, they have succeeded. (And) for financially motivated actors, their goals are centered on making money."

Users encounter a wide variety of these exploits. Even when the end result looks similar - like a pop-up that congratulates a user on winning a gift card - the underlying code shifts like a swiftly evolving metavirus. Once a bad ad is detected and blocked, exploiters, hackers and cyber criminals must find new ways to cloak the bad ad. Think of it as wearing a fake beard to foil facial detection software.

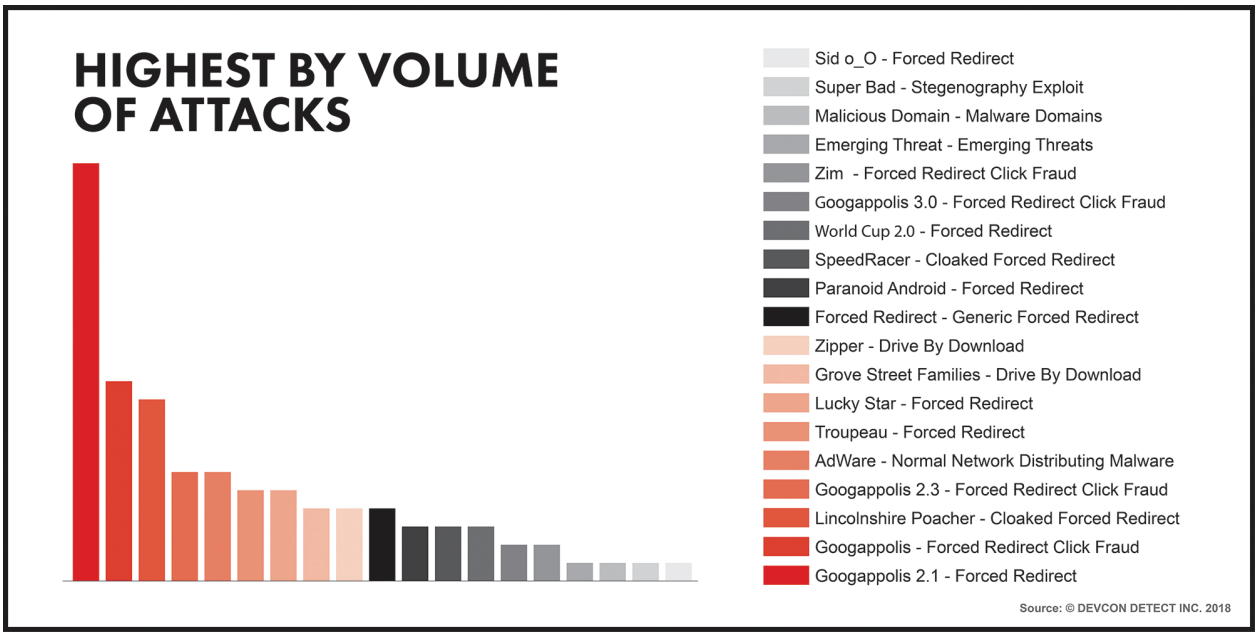

(Note: The names of exploits on the chart are arbitrary. The researchers and developers at DEVCON involved in the study "named" exploits as they were discovered to help flag and sort them into categories. Hence the sometimes colorful or humorous names. They were not named by those who caused the exploits.)

"We identify unique code signatures within the execution of an exploit and use the DEVCON platform to block bad ads based on those signatures. But the work doesn't stop there," said Christina Brown, Hackmaster General at DEVCON. "Cyber criminals are constantly changing those signatures, requiring new signatures to be added to the platform regularly. When massive attacks are deployed at one time, it can feel a little bit like whack-a-mole."

A known trait of online advertising – and one of the major drivers of the rise of programmatic networks – is the imbalance between supply and demand. That is, the supply of sold advertising – locally, regionally and nationally – and the amount of audience demand outpacing that supply. So, during major news events or stories that go viral (Think: "Florida Man...") there is far more audience that can keep pace with higher-CPM advertising – and that local advertising would be wasted on out-of-market users.

To make money off these spikes in traffic, digital publishers tap into programmatic networks by setting floor rates for filler ads. The lower the floor price, generally, the more generic the ad. While well known in the industry, most online users are unaware of this practice – or how this can contribute to exploits entering the ecosystem.

Ad networks have many protections against these bad actors dirtying up their supply pipes. They verify and whitelist suppliers and publishers.

"There are dozens of safeguards that we deploy to stop bad actors. Our team is constantly scanning tags and sites for issues," said James Byrd, COO of OpsCo, an advertising and yield operations resource for publishers. "We prescreen new partners for a full week before pushing them live, setup test pages for each demand partner to isolate testing and run reports on buyer activity to check for unusual spikes that would point to a possible exploit.

"Ensuring quality demand for our supply is a full time job. "

Yet when floors fall – to, say, $0.35 per 1,000 impressions – the chances of bottom-feeding and even fake networks getting into the mix increases.

But in the AdHack.org study, researchers also found exploits targeting high-traffic events – but not by ad impressions. They targeted by type of content.

During the 2018 midterm elections, the AdHack.org study showed just how targeted exploits could be. Exploits spiked 20,000 percent over norms (from .45 percent to 2.5 percent) for local news consumers seeking election return updates and news stories (and URLs) containing political terms such as those listed to the right.

TARGETED TERMS

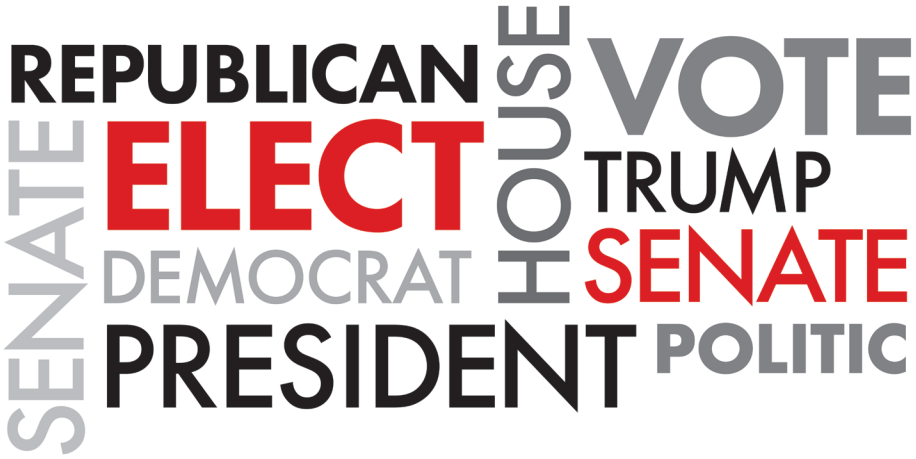

Most affected was Apple's iOS platform, at 94 percent. All other operating systems – Windows desktop, Android and Mac desktop – filled the remaining gap.

At first blush, researchers thought a spike in interest and traffic to news sites would account for the aggressive increase in exploits. But the upcoming shopping season that kicked off with the night before Black Friday through Cyber Monday disproved that traffic was corollary to the number of bad ads coursing through local media.

As the following chart indicates, while the number of ad impressions on local news sites remained largely consistent from Thanksgiving through Cyber Monday, the number of ads being served that contained exploits began accelerating the night of Thanksgiving, with Black Friday (11/23/18) showing a 267 percent increase in exploits and climbing to its peak of a 528 percent increase on the Sunday (11/25/18) before Cyber Monday.

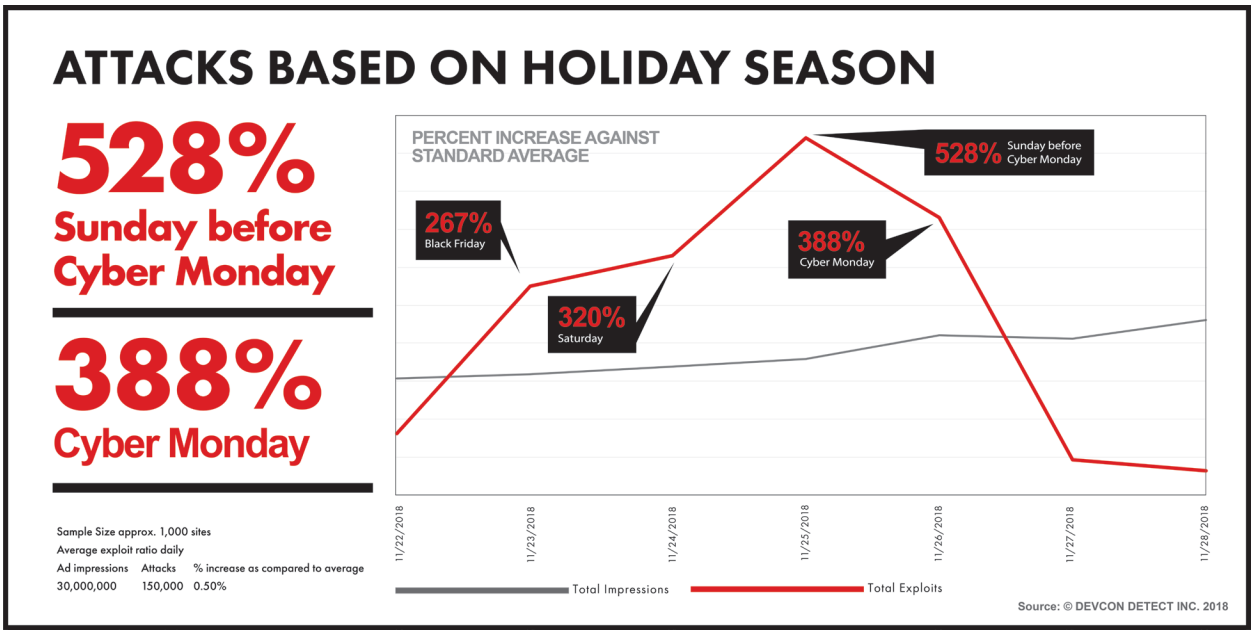

It's easy to understand they psychology behind this dramatic climb. Online consumers, already looking for bargains and sales following Thanksgiving, were easy targets for exploits promising $10,000 gift cards from Amazon or Walmart – while first being sent through a cascade of forced-redirect sites to drive up ad impression traffic on faux news and shopping sites.

In fact, it opens a window into how criminals behave.

FROM BANNER ADS TO THE DARK WEB

Many executives with digital publishing and revenue responsibilities aren't aware of how vulnerable advertising systems can be.

"The ad ecosystem is very appealing to cyber criminals because it is completely unregulated with a ton of access to money and data," Anaya said. "That large amount of money combined with billions of users worldwide gaining more and more access to digital media – which increases their exposure to digital ads – make the ad ecosystem very lucrative."

What's also very appealing is the lack of regulation. The lack of oversight and governance translates to the reduced risk of being exposed and getting caught. The ad ecosystem offers bad actors levels of anonymity rarely seen elsewhere.

And they collaborate.

"Cyber threat actors do not work in a vacuum, they collaborate with each other," DEVCON's Anaya said. "They find others on the Dark Web that could enhance their own capabilities."

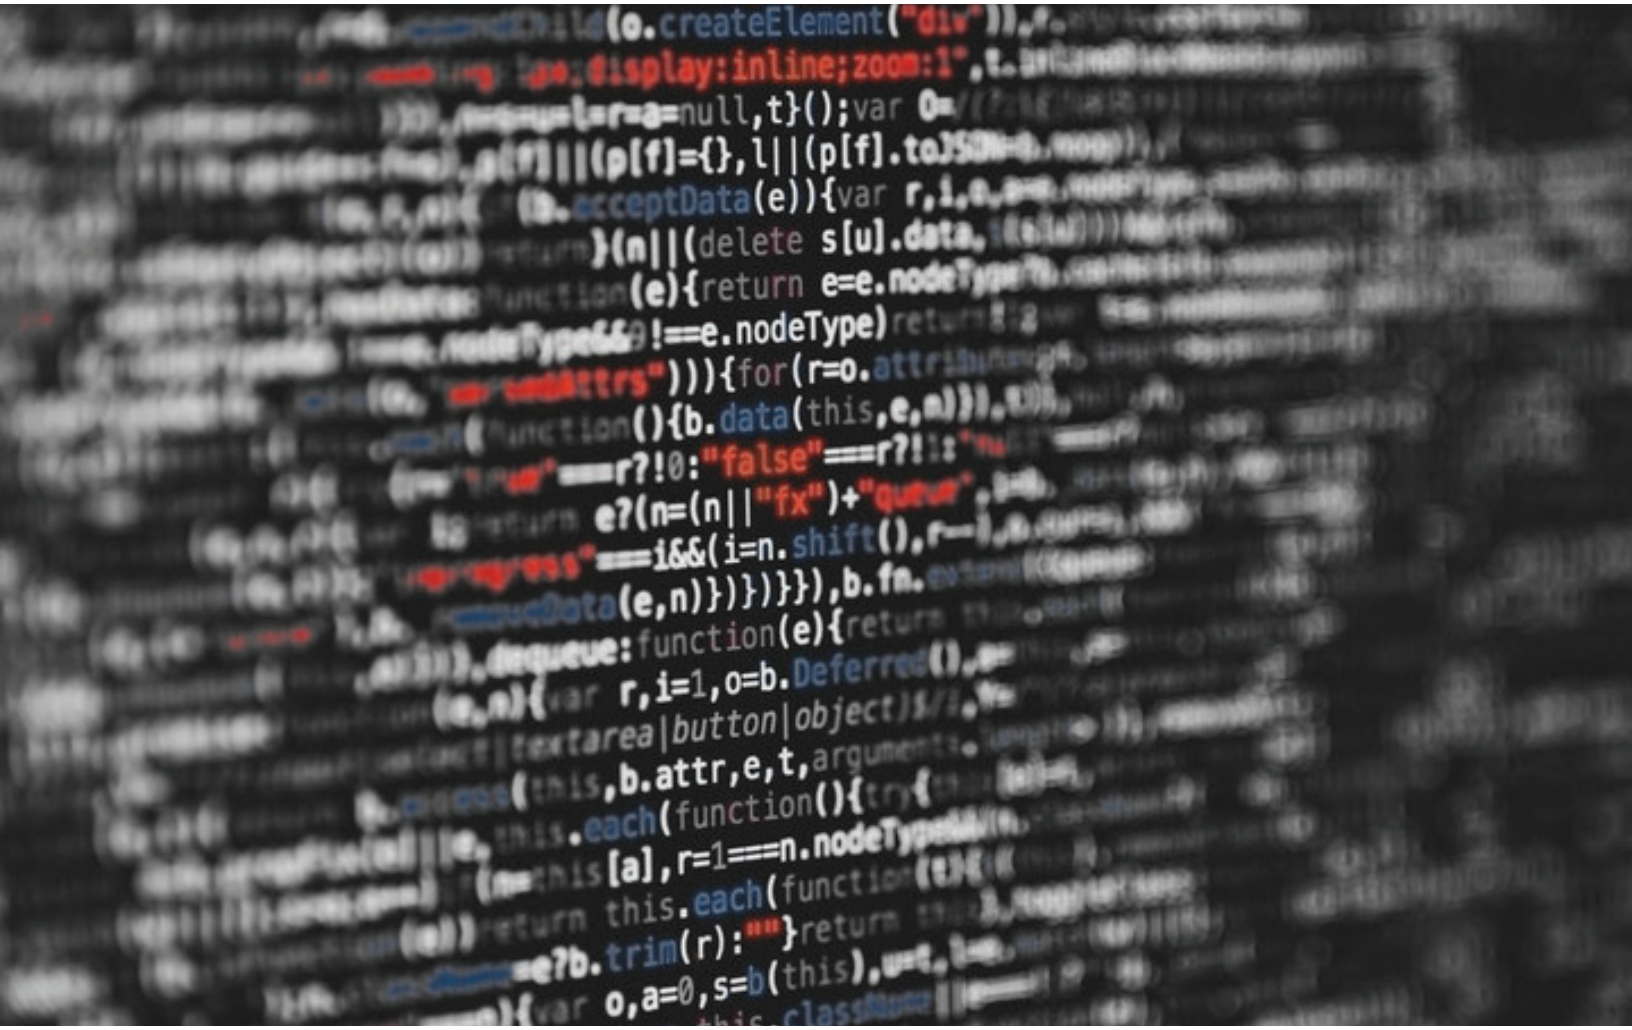

In other words, the bad actors are much more organized and collaborative than publishers, ad networks, ad operations, and everyone else on the "good" side.

While "the Dark Web" may sound cinematic, reality is simple and transactional, but populated with criminals and would-be criminals. The Dark Web is a collection of websites that exist on an encrypted network such as Tor and cannot be visited by using traditional browsers or found by using standard search engines. In order to access the Dark Web, users have to use a browser built to access it. Once someone has accessed the Dark Web, they could visit a vast array of nefarious marketplaces.

"The vast majority of the Dark Web is used for illicit purposes, such as the distribution of malicious software (malware), drugs or child pornography," Anaya said. "The Dark Web is where cyber threat actors buy and sell goods and services, build relationships, and learn of new and innovative ways to exploit victims."

Ad exploits are rarely the work of a single bad actor who builds an exploit, uses it, and collects all the funds from its use. Instead, each bad ad is part of a familiar supply-demand chain. Someone creates it, he/she then sells it to a number of buyers, the buyers use it and then the buyer may use another group to collect the funds, depending on the scheme. Given the nature of the Dark Web and crypto-currency, each person or group might know of another one, but not in true name.

And since all the exchanges are typically done using encrypted and untraceable channels, it's nearly impossible for law enforcement to track these transactions.

Which means that it's up to local news sites and the networks supplying their programmatic ads to self-police. And that's not just good housekeeping or (yet another) unbudgeted expense.

Fighting ad fraud drastically improves the bottom line.

WHY IT'S MORE PROFITABLE TO FIGHT AD FRAUD

On the front lines, the publisher's warchest to fight fraud is limited in either resources or knowledge. In the majority of local publications and broadcast outlets across the country it's usually both. Because of these limited resources, the local battle against ad fraud always has a direct impact to their digital revenue potential.

"On a non-technology level, publishers really only have two ways to fight malicious exploits that plague their users, 'Turn If Off' or 'Hedge'," said OpsCo's Byrd. "This boils down to the publisher giving up revenue in turn for 'solving' the problem."

A local site "turning off" programmatic can be as little as pausing individual networks to putting a complete stop to the revenue pipeline. Most local news companies have deployed this method out of sheer desperation to solve the onslaught of external (and internal) complaints. This method relieves the outside pressure but can't sustain the business pressure of needing the revenue flowing. At some point, a publisher has to turn things back on.

The "on/off' tactic, often handed down from the executive level to the operational level, results in the second way Byrd says publishers can fight ad fraud, 'Hedging."

When the extremely reactive approach of "turn it off'"becomes a burden on internal operations, the folks with their hands on the controls do what they do best: try to solve the problem. "Hedging" is a direct result of someone with a little knowledge of the programmatic operational space attempting to solve the issues that are ultimately making their work harder and the overall revenue sink.

"Hedging is the yield manager pulling all the levers at their disposal to stop the issues," Byrd said. "And the chief way that publishers 'hedge' against ad fraud is floor prices."

Raising mobile floor prices cuts out the lower-price band reducing the overall reports of issues. But in reality, all this is doing is cutting down the inventory available for monetization, thus reducing the number of issues. Publisher are simply not monetizing as many impressions.

In 2018, Byrd said his team at OpsCo heavily relied on "hedging" to protect all the publishers that rely on them for monetization. Over the course of that year they reported more than 1.9 billion programmatic mobile impressions across their sites. Out of all those impressions 'hedging' resulted in 20.21% of those impressions being sacrificed to help stop malicious exploits.

OpsCo's attempt at fighting ad fraud themselves was not cheap, says James. "We missed out on just under $250,000 in mobile revenue for our pubs that year."

How much have you lost fighting ad fraud with the binary On/Off and Hedge options? With local sites being strapped for resources and knowledge this question is sometimes impossible to answer. To help address this gap, and as a result of this study, AdHack.org and DEVCON developed a calculator for local sites or media companies to shed light on how much is at stake. It's at **AdHack.org/calculator**.

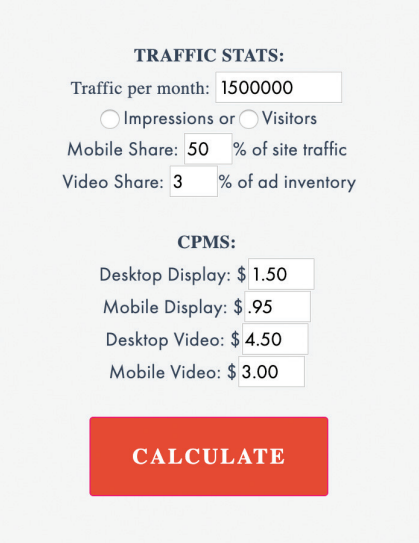

But sometimes the motives of the bad actors aren't financially based, as the AdHack.org study detected in early 2019.

NATION STATE TRACKING PIXELS

In February, the DEVCON detection tools deployed by participants in the AdHack.org study discovered Russian tracking pixels showing up in the ad tags and creative on local media sites. That list quickly expanded to include tracking pixels from China, and more than a dozen others.

Tracking pixels can be used to target attacks based on gender, location, income, job title, political affiliation and hundreds of other demographics. In authoritarian countries, data gathered by ISPs and apps is routinely shared with the government as a way to monitor and control citizens.

So what are they doing showing up on a Midwestern college-town sports page?

How they got there has a variety of answers, some benign and others more sinister. On the benign side, to fulfill ad demand from programmatic networks, media companies open bidding to hundreds or even thousands of suppliers, and not all of them are domestic – or are served from computer networks outside the United States.

On the more sinister side, DEVCON's Anaya said nation state actors also buy exploits on the Dark Web to obfuscate their activity and minimize anyone's ability to attribute the action to them.

"In the past, when a nation state group would generate and deploy a custom piece of malware, it served as a way for the victim or their government to trace the origins of the exploit back to the county in question," Anaya said. "Nation state actors evolved and began to purchase common exploits to cover their tracks."

Once one of these pixels is introduced to a user, any number of demographics, behaviors and political affiliations can be tracked. This information can be used to target ads but also content, as seen in the 2016 elections.

Which means bad actors aren't just defrauding publishers of billions of dollars or exposing their users to harmful malware. The bad actors are using ad fraud to weaponize local news sites at a clip of millions of impressions a day.

Trackers from Russia on local news Sites

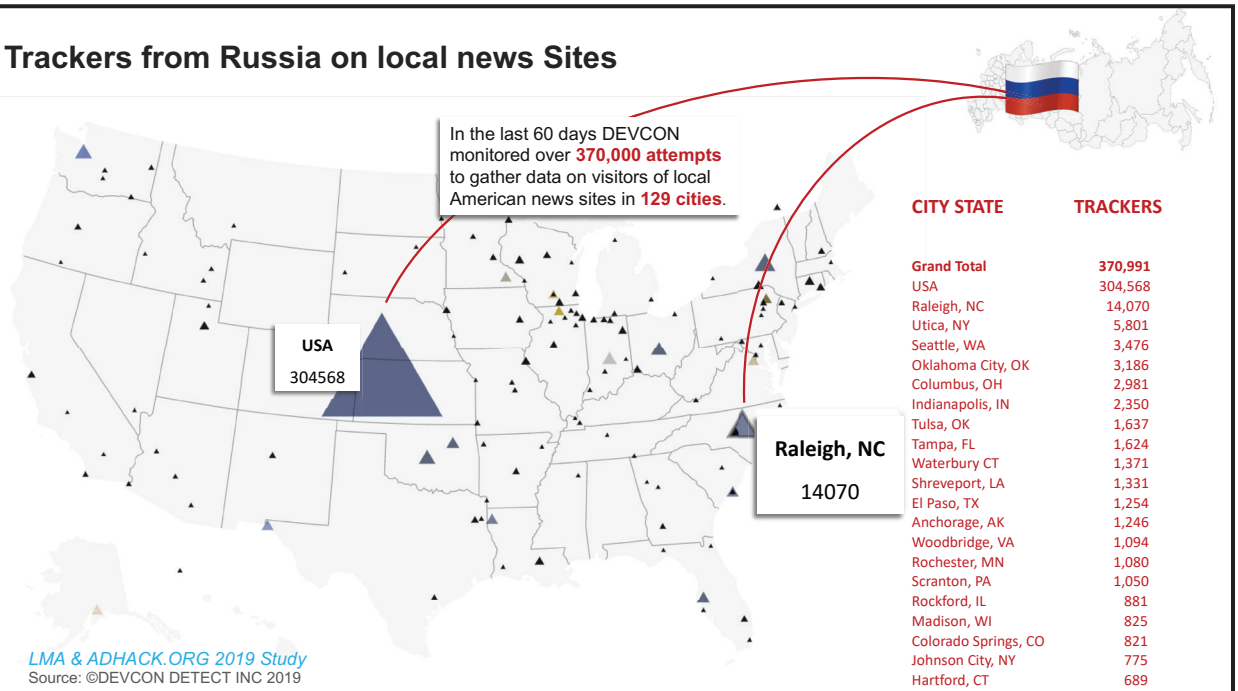

CITY STATE	TRACKERS
Grand Total	**370,991**
USA	304,568
Raleigh, NC	14,070
Utica, NY	5,801
Seattle, WA	3,476
Oklahoma City, OK	3,186
Columbus, OH	2,981
Indianapolis, IN	2,350
Tulsa, OK	1,637
Tampa, FL	1,624
Waterbury CT	1,371
Shreveport, LA	1,331
El Paso, TX	1,254
Anchorage, AK	1,246
Woodbridge, VA	1,094
Rochester, MN	1,080
Scranton, PA	1,050
Rockford, IL	881
Madison, WI	825
Colorado Springs, CO	821
Johnson City, NY	775
Hartford, CT	689
Salt Lake City, UT	610

Trackers from Nigeria on local news Sites

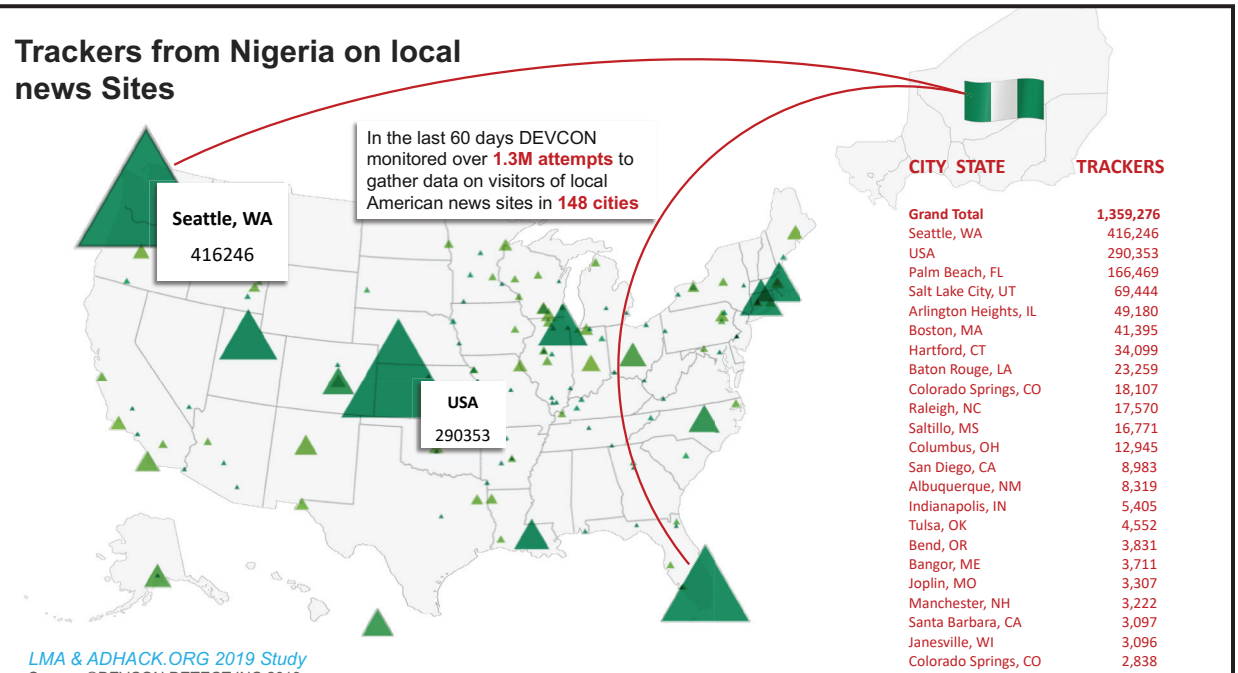

CITY STATE	TRACKERS
Grand Total	**1,359,276**
Seattle, WA	416,246
USA	290,353
Palm Beach, FL	166,469
Salt Lake City, UT	69,444
Arlington Heights, IL	49,180
Boston, MA	41,395
Hartford, CT	34,099
Baton Rouge, LA	23,259
Colorado Springs, CO	18,107
Raleigh, NC	17,570
Saltillo, MS	16,771
Columbus, OH	12,945
San Diego, CA	8,983
Albuquerque, NM	8,319
Indianapolis, IN	5,405
Tulsa, OK	4,552
Bend, OR	3,831
Bangor, ME	3,711
Joplin, MO	3,307
Manchester, NH	3,222
Santa Barbara, CA	3,097
Janesville, WI	3,096
Colorado Springs, CO	2,838
Oklahoma City, OK	2,370

CONCLUSION AND NEXT STEPS

Thanks to its partnership with the Local Media Association, DEVCON and OpsCo, the ongoing AdHack.org study is moving the conversation of digital ad fraud in local media from episodic incidences to a more universal view. That view is frightening in its scale and complexity, but encouraging in its outlook. For this first phase of the study, our key takeaways are equally cautionary and optimistic:

- Local news consumers are exposed to thousands of bad ads containing malware, ransomware, forced redirects and numerous other exploits each day.
- In fact, in the 20 minutes it took to read this white paper, more than 487,000 exploits were detected in the AdHack.org test group.
- Exploits find their way into ad networks by criminals spoofing the system for personal monetary gain or to inform future exploits, including targeting content based on behavior and political affiliation, or to inform a "watering hole" ransomware exploit on someone based on job title or household income.
- While publishers and ad networks take numerous precautions to prevent exploits from slipping through, bad actors are well informed and networked via the encrypted "Dark Web," and exploits continually evolve – often at a rate of once a day.
- Trackers from servers in more than a dozen nation states have been detected in local news programmatic ads.
- Typically, less than 1 percent of all ads served – 0.45 percent – are detected to be malicious. While that seems small, it equates to millions of news consumers exposed each month.
- Ad fraud robs local news publishers and broadcasters of revenue while pushing away news consumers who are exposed to exploits and malware.
- Fraud detection for publishers often is considered an unbudgeted expense – but in fact can drive significantly more revenue than either ignoring the problem of following the "Shut it Off/Hedge" options.
- Fighting ad fraud must be a collective effort between publishers, networks, agencies and brands. No one entity is solely responsible – but it is consumers who are most at risk.
- As bad actors and their exploits continue to evolve, so will the need for more sophisticated detection technology – and further study and discussion.

"What keeps me up at night is the complete lack of barriers to entry. I've never seen anything like it. It's really only limited to the imagination of the attacker to what exploits they can distribute through the online advertising system - everything from consumer data harvesting scams to remote access trojans."

JOSH SUMMITT | CTO & COFOUNDER | DEVCON

ACKNOWLEDGEMENTS

This phase of the AdHack.org study would not have been possible without the following participants and their generous contributions to illuminating the issue of digital ad fraud for local news sites and their consumers:

- Local Media Association: Local Media Association (LMA) serves more than 3,000 newspapers, TV stations, digital news sites, radio stations, directories and research & development partners. It is the only industry trade organization that brings all local media together for the purpose of sharing, networking, collaborating and learning. Contact Jed Williams, Chief Innovation Officer, Jed.Williams@localmedia.org
- DEVCON: DEVCON is committed to building tools that preserve the media industry and gathering intelligence that tracks down and stops cyber criminals. Contact: Casey Hester, COO, chester@devcondetect.com
- OpsCo: OpsCo is a digital-publisher shared services agency specializing in ads, yield and data. Contact: James Byrd, COO, james@ops.co
- AdHack.org is a nonprofit dedicated to raising awareness of the costs of ad fraud for news publishers and consumers. Contact: info@adhack.org