THE CHANGING FACE OF MALVERTISING

Evolving tactics blur the lines between bad ads, scams & overt malware

THE MEDIA TRUST
We know digital security.

State of Malvertising

Malvertising incidents continue to rise despite widespread adoption of blocking techniques

When it comes to malvertising, the more things change the more they stay the same. In 2020, The Media Trust detected an unprecedented rise in fraud, malware, and disinformation on the web—beginning with the COVID-19 pandemic, use of celebrity-laden click bait, continuing through global economic uncertainty, and culminating with the U.S. national elections. Digital advertising is responsible for the propagation of these consumer-directed threats.

In 2020 The Media Trust detected more than 20,000 distinct malvertising attacks and blocked billions of bad ad impressions—a 5X block increase compared to 2019. Typically, these distinct events penetrate dozens of publishers and AdTech partners, poisoning the user experience or infecting millions of consumer devices at a time. With a 724% incident increase since 2014, malvertising continues to plague the industry despite the introduction of new mitigation tools and techniques. [Figure 1]

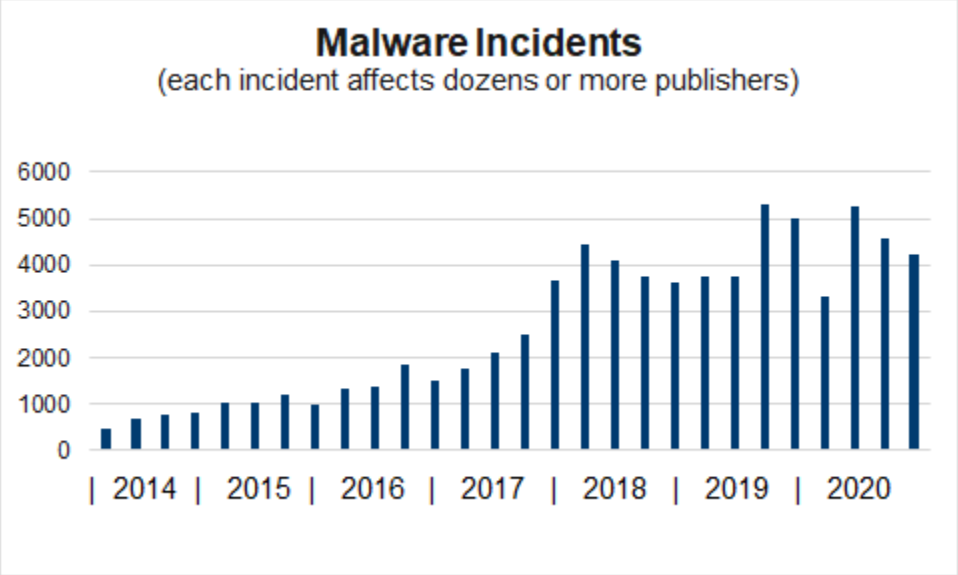

Figure 1: Number of distinct malware incidents over the past six years. Each incident typically affects dozens of Publishers and AdTech providers.

Emerging Malvertising Trends

Bad actors increased use of basic tactics while probing new sophisticated strategies

Analyzing the malvertising incidents of 2020 reveals a shift in the mindset of bad actors. While they continue to look for the most efficient method for deploying their attacks, the tactics have changed. Bad actors are using more scam-oriented campaigns—not overt malware or redirects—and adopting new strategies of varying levels of sophistication. This shift was further accelerated by the COVID-19 pandemic.

Key malvertising trends include:

- **Resurgence in clickbait creative**: As a means of attracting user action, celebrity and "too good to be true" promotions increased dramatically—1500% throughout the year. The pandemic also encouraged campaigns promising hard-to-find items including face masks and sanitizers, which upon investigation revealed goods of dubious integrity. At one point these scams were 27X more frequent than usual.
- **Rise in inflammatory creative**: Due to the uncertainty driven by the pandemic, state of the global economy, social justice movements, and U.S. elections, consumers are much more sensitive to inflammatory content. Increasingly, consumers believe this kind of advertising content is more harmful than overt malware and redirects due to its ability to evoke a negative emotional reaction. [Figure 2] Detection of inflammatory ads almost doubled, with political-oriented creative contributing to more than 75% of the category's violation criteria each month throughout the year. Even worse, bad actors adopt inflammatory creative and target it to consumers most likely to be receptive to the message.

Figure 2: Sampling of inflammatory creative found on premium publisher sites

- **Longer, more frequent test phases**: Large-scale attacks are now typically prefaced with several test phases that execute for a few days over the course of several weeks, sometimes months. While the patterns and objectives do not change, the test attacks rotate through different AdTech providers, probing the ability to bypass detection defenses. If not remedied at the buyer level, the attacks launch and affect dozens of premium publishers.
- **Targeting of "at-home" environments:** Malware targeting traditional non-corporate environments—ISPs, mobile devices, and mobile networks—surged 20X at the pandemic's outset. While the malware infection rate for consumers accessing the Alexa 1,000 fell in the summer, it remains at an elevated rate—an average 2 times per 1,000 visitors, which is more than double 2019's typical rate. If every mobile user accessed one website a day, there would be at least 2 billion devices exposed to malware each day.
- **Blocker and ad server evasion**: Bad actors are now regularly identifying and bypassing creative checkers and malware blocking tools. Routinely detecting active malvertising incidents on websites using these tools, The Media Trust has observed these bad actors learning how to bypass detection by adopting new techniques. As malicious actors find new and creative ways to circumvent existing safety measures, it is more important than ever to combine real-time blocking with sophisticated tag and site-scanning analysis to capture emerging threats.
- **Enhanced cloaking efforts**: A key change is the more regular and sophisticated use of cloaking, a tactic to hide malware and evade detection by delivering different versions of the campaign creative based on the user's profile. Bad actors are combining hyper-obfuscation (often 2,000+ lines of code), standard library tweaks, and spoofed domains to deliver different

creative based on the user's geography and device. Depending on the user's profile, the creative delivers extra malicious code. The Media Trust is actively working with TAG Threat Exchange to block these threats in real time to minimize the attack scale. The group is collaborating on best practices to identify and terminate these bad actors across the industry's largest DSPs and SSPs.

- **Landing page-oriented compromises:** Bad actors have discovered that malware blocking tools don't catch malware on landing pages. Typically, 10% of all malware detected, malvertising found only on landing pages has jumped to more than 18% of all incidents. As the year went on celebrity scams changed from redirects to compromised landing pages as their primary vector. When combined with cloaking, it can be very pernicious.
- **Ubiquity of fingerprinting:** What was once an occasional sighting, fingerprinting code is regularly detected in scans. Throughout the year fingerprinting code was observed to be identifying devices, laying the groundwork for future device- and platform-oriented attacks.

On average, approximately 90% of client-side executing code is provided by third-party partners (Digital 3PC). Because technology teams only manage 10% of their digital asset code, bad actors take advantage of this unmanaged attack surface to penetrate websites and/or mobile apps and cause harm. Compromise of third-party partners, like those in the advertising ecosystem, is now a primary attack vector. This report focuses on advertising-delivered malware.

Top Named Digital Threats

Most significant incidents evade detection and blocking defenses to target consumers

In 2018, our Digital Security & Operations team (DSO) began to name significant threats that exhibit unique characteristics to avoid detection or quickly scale to affect large groups of clients. Requiring extensive analysis, these incidents are usually grouped and named, with the name summarizing the threat actor's primary characteristic or leveraged vulnerability. Each named threat can affect not only dozens of AdTech and publishers but also millions of impressions.

Most of these named incidents are appended with "3PC" to represent the use of third-party code in the malware's delivery, which is the main source of digital malware propagation. To date, DSO is tracking more than 20 active named digital threats.

An **Incident** is a unique piece of flagged or unnecessary content with matching characteristics according to subdomain, domain, IP, creative, referral source and behavior, e.g., redirect, exploit kit, toolbar.

A **Named Threat** is a group of incidents causing significant harm across the broader digital advertising ecosystem.

This report will focus on the top 5 causing the most trouble across the global digital ecosystem in 2020.

CELEBRITY SCAMS

First detected in July 2019, Celebrity campaigns (also known as FizzCore) exploit thousands of victims by manipulating ad tech to drive clicks and redirects to cryptocurrency investment scams. Users are targeted with “celebrity-endorsed” click bait that leads to a landing page promoting a cryptocurrency scheme.

These scams implement evasion techniques—such as cloaking—to elude domain and pattern-based ad quality reviews. The malware uses two cloaked scripts: one for targeted and another for non-targeted users. If the device meets the targeted geographic, device or behavior parameters, the cloaked script loads the preliminary malicious landing page, as well the cloaked ad creative. To date, the malware attack has been detected across several, well-known DSPs and SSPs.

Consumer Harm:
- Financial scam/loss
- Misinformation propagation

Defend against Celebrity scams:
- Continuous, real-time scanning
- Blocking tool that uses a frequently updated exclusion list

Celebrity Scams stepped up their attacks in 2020, driving a 15X incident increase since we changed the classification from scam to malware. [Figure 3] Toward the beginning of 2020, the rate of affected users outpaced the number of distinct outbreaks. This is primarily because programmatic price floors dropped in the wake of a global ad spend slowdown, providing bad actors easier access to the digital advertising ecosystem.

The increase in the latter part of the year significantly affected European publishers and AdTech partners. As most detection is performed via US-based scanners, bad actors are more successfully targeting Europe-based profiles to evade detection and blocking.

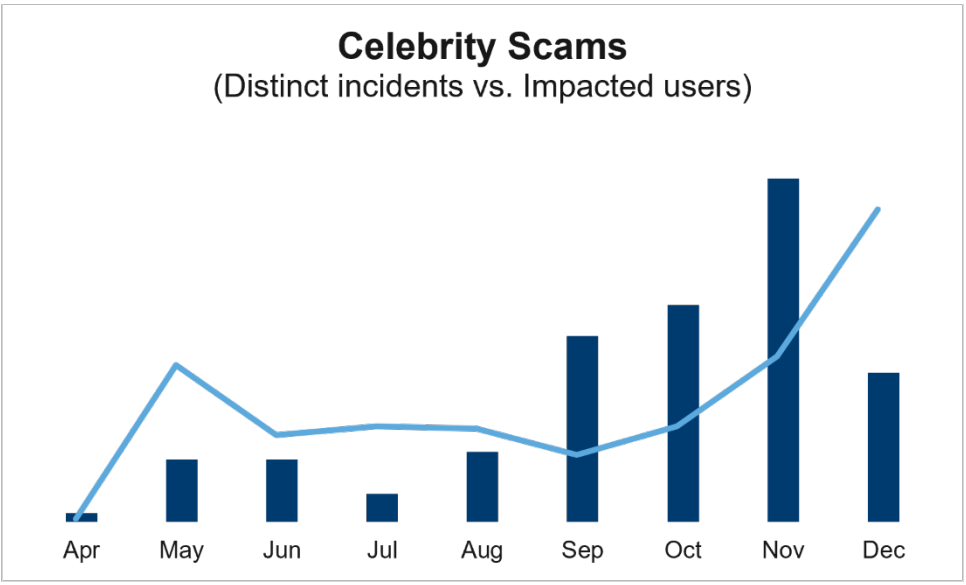

Figure 3: Number of distinct “Celebrity Scam” incidents or outbreaks and number of users impacted.

The scam evolved throughout the year. Initially, the scam used free and open-source ad servers to serve celebrity creative to targeted users in Europe, with an emphasis on the UK and Germany. As the

year progressed, the scams moved to commercial ad servers to deliver broader clickbait-oriented creative (e.g., “get rich quick”) to users in dozens of countries across six continents. In addition, the creative has morphed to a more native style that is served programmatically, and the landing pages may feature other dubious investment schemes, not necessarily bitcoin.

CORONAVIRUS

The COVID-19 pandemic drove a significant increase in scams where Coronavirus-related campaigns promoted goods and/or services with misleading or false claims related to the pandemic—i.e., ineffective face masks, hand sanitizer, unverified cures, etc. The campaigns began in earnest February 2020 with the intent to extract payment from unsuspecting users by offering illegitimate goods or in-demand supplies at inflated prices. After interacting with these domains, users could be targeted by phishing emails from fraudsters in an attempt either to plant malware on their computer or extract personal information.

Consumer Harm:

- Poor user experience due to scam ads
- Fraudulent purchases
- Phishing
- Misinformation propagation

Defend against Coronavirus scams

- Continuous, real-time scanning

Coronavirus scams reached a noticeable volume in March, leading to a standalone malware classification. Not surprisingly, almost every incident affected North American users (88%), with European users affected 4% of the time. [Figure 4] Peaking at 200X in May, the campaigns slowly tapered off throughout the year. But they remain and will continue in 2021, likely with vaccine-related messaging.

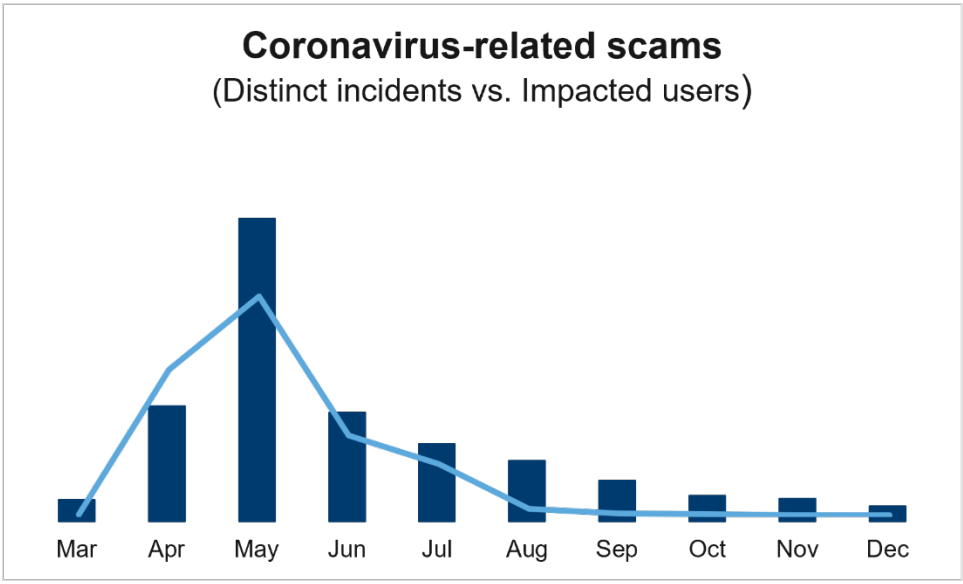

Figure 4: Number of distinct “Coronavirus-related Scam” incidents or outbreaks and number of users impacted.

GHOSTCAT-3PC

This malicious campaign, initially detected in June 2019, uses advanced obfuscated code and delivery patterns to evade signature-based defenses often used by publishers. The malware powering the attack bypasses conventional blockers to hijack mobile browser sessions, primarily those in the U.S. and Europe.

Once the malware commandeers the user's browser session, it disables settings that block unwanted advertisements and delivers adware to the device. On Android devices version 5.1 or earlier, GhostCat attempts to install the Android.Xiny.5260 trojan, which facilitates the download of additional malicious apps.

GhostCat continued to regularly target U.S. and European users throughout 2020, with these users experiencing 82% of all incidents. Concatenation became more obfuscated, frequently adopting techniques and coding patterns used by other known actors such as ICEPick-3PC. The attack still only executes on mobile, primarily Android; it has not been detected in desktop environments to date.

Consumer Harm

- Unwanted programs in the form of adware
- Auto-download of malicious apps
- Redirects to fraudulent pages (e.g., fake gift cards)

Defend against GhostCat-3PC

- Continuous, real-time scanning

GhostCat is another example where the rate of affected individuals outpaces the growth in the number of incidents. [Figure 5] Individual infections peaked in March and April before leveling off for most of the remainder of the year. This is because there were several more incidents (based on different creative, referrer domain or malicious domain) that affected relatively less users in June and July. The rapid pace of creative rotation and payload domains forces a shorter incident lifespan, thereby affecting less users per incident. However, these maneuvers were still detected, so the actor reverted to previous tactics and, as a result, incident count and users affected realigned in the fall.

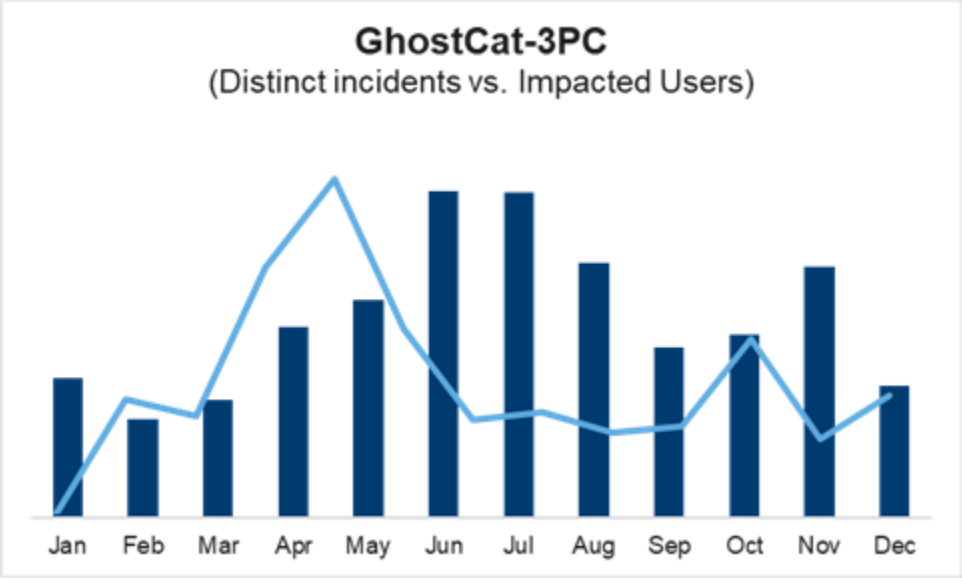

Figure 5: Number of distinct "GhostCat" incidents or outbreaks and number of users impacted.

ICEPICK

First detected in October 2018 and sometimes referred to as eGobbler, the ICEPick-3PC campaign leverages compromised third-party tools—used to implement interactive web content and animation in digital advertisements, e.g., HTML5—to redirect and exfiltrate sensitive user and device information. These third-party tools are often pre-loaded onto client platforms by self-service agencies.

Utilizing built-in JavaScript functionality permitted by the browser, the malware uses the WebRTC platform to establish a connection between the infected device and a remote host. This connection harvests device-specific data including device model, browser version, and IP address. Usually targeting vulnerable mobile Android devices, the final stage of the attack initiates a prompt for the user to download a fraudulent Google Play Store application, which delivers persistent adware to the device.

Consumer Harm

- Redirects
- Phishing
- Unwanted programs in the form of adware

Defend against ICEPick-3PC

- Continuous, real-time scanning

Since coming on the scene in 2018, ICEPick has remained a significant malvertising actor, tripling its presence across the digital advertising ecosystem compared to 2019. [Figure 6] Incidents primarily affected users in North America (47%) and Europe (35%).

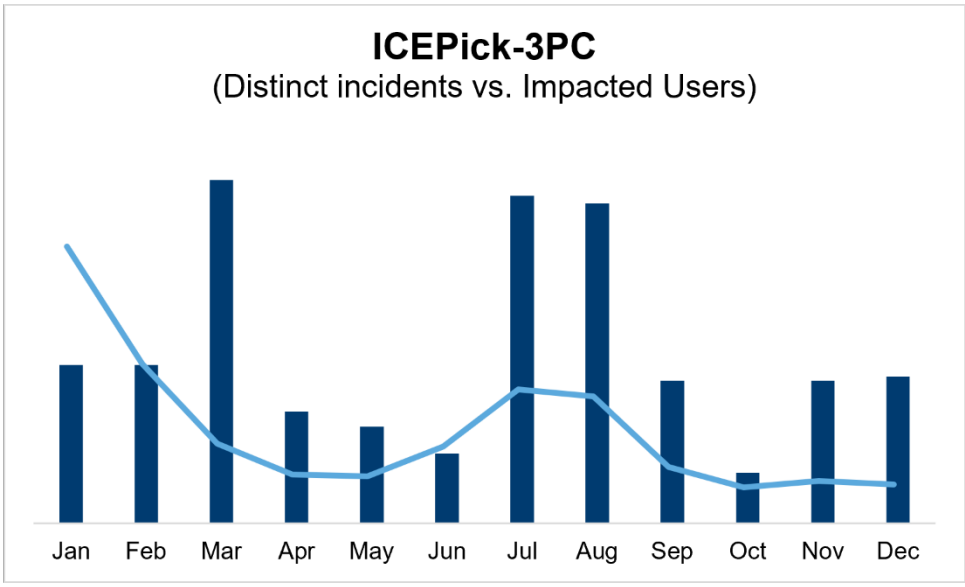

Figure 6: Number of distinct "ICEPick" incidents or outbreaks and number of users impacted.

In 2020 scammers adjusted tactics. In addition to expanding into desktop environments, the actor launched several small campaigns testing different platforms and domain pairs. From January through March, dozens of campaign flights with a few tags affecting a handful of smaller publishers started on Thursdays and continued through Sundays. This was in preparation for a larger-scale attack that occurred over three successive weekends during July and August.

LNKR

This malicious campaign targets and tracks users by injecting JavaScript code via compromised Google Chrome browser extensions to add malicious content directly on a page during runtime. A user device with the LNKR adware will have ads injected in odd places that often break the page. The malicious code delivers illegitimate ads for monetization, redirects users to malicious content, and tracks browser activity. LNKR also has the ability to search for pages with user write access, which allows it to upload other malicious JavaScript in addition to its own.

Consumer Harm
- Phishing
- Redirects
- Unwanted programs in the form of adware

Defend against LNKR
- Real-time, continuous scanning
- Client-side blocking

First named in early 2019, LNKR detection in the broader digital advertising ecosystem was sporadic. In late 2020 LNKR infections reached a significant number to warrant threat group tracking. [Figure 7] To date, the campaigns are primarily affecting users in North America (42%), Europe (21%), and Asia (17%).

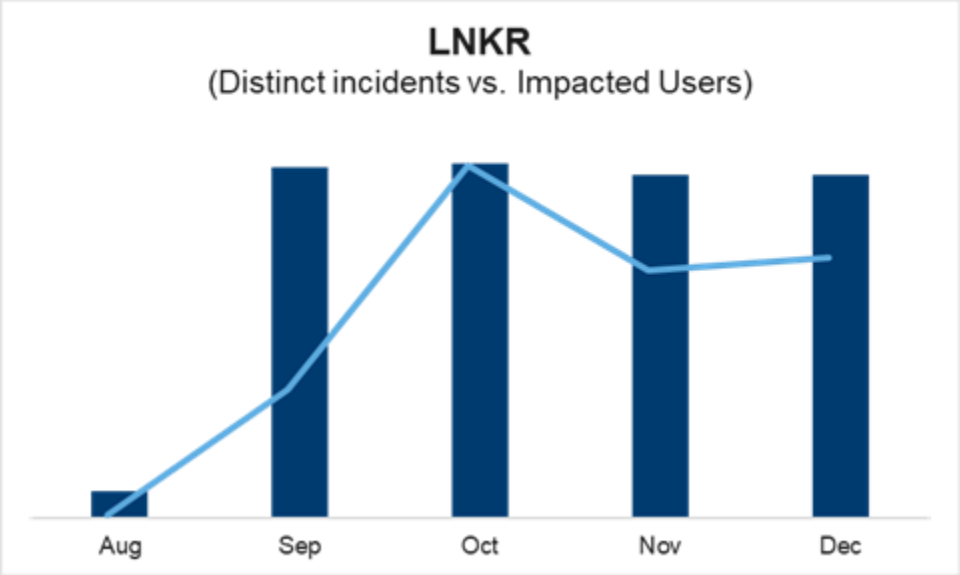

Figure 7: Number of distinct "LNKR" incidents or outbreaks and number of users impacted.

LOOKING FORWARD

Malvertising is rapidly changing. To keep abreast of its evolving nature, The Media Trust uses several techniques including:

- Increased deployment of more random user profiles to capture emerging threats
- Landing page click-throughs to analyze the full user experience—creative, tag, and landing page—and identify scams
- Regular update of our exclusion list to block threats as they emerge

- Enabling more “at-home scanning” profiles. These profiles currently represent 10% of our scanning and there are plans to double it throughout 2021
- Testing of various fingerprinting techniques to elicit malware and unauthorized tracking

In addition to our close relationship with Google and multiple enforcement authorities, The Media Trust is the inaugural malware vendor for the TAG Certified Against Malware seal and regularly dissects and presents evolving threats to the TAG Threat Exchange. These relationships not only help us collaboratively terminate bad actor access to the digital ecosystem, but also support prosecution and other deterrence efforts.

For 2021, publishers and AdTech providers need to:

- Consider how increased “at-home” access will affect their users
- Review changing geographic patterns—there are fewer metropolitan users—and verify these user experiences can be safeguarded
- Update exclusion lists—ideally multiple times a day—with original-source malware data instead of third-party data, which drive false positives
- Keep abreast of fingerprinting and how it will drive more cloaking tactics
- Be on the lookout for continued scams—e.g., celebrity clickbait, Coronavirus vaccine, etc.
- Evaluate malware risks from non-advertising sources—e.g., payment card theft via compromised web code on subscription pages
- Assess partner relationships with an eye to those responsible for delivering poor ad quality

About The Media Trust

The Media Trust is on a mission to fix the digital ecosystem. Through continuous monitoring of websites and mobile apps, we provide transparency into the complex relationships delivering the consumer experience. More than 600 premium enterprises, media publishers, ad networks/exchanges, and agencies—including 40 of comScore's AdFocus Top 50 websites—rely on The Media Trust to identify and remediate security, data protection, and quality risks that can lead to regulatory fines, depressed inventory value, revenue loss, and brand damage. For more information, visit www.mediatrust.com.