On the Infrastructure Providers that Support Misinformation Websites

Catherine Han
cathan@stanford.edu
Stanford University

Deepak Kumar
kumarde@stanford.edu
Stanford University

Zakir Durumeric
zakird@stanford.edu
Stanford University

ABSTRACT

In this paper, we analyze the service providers that power 440 misinformation and hate sites, including hosting platforms, domain registrars, DDoS protection companies, advertising networks, donation processors, and e-mail providers. We find that several providers are disproportionately responsible for hosting misinformation websites, most prominently Cloudflare. We further show that misinformation sites also disproportionately rely on several popular ad networks and payment processors, including RevContent and Google DoubleClick. When misinformation websites are deplatformed by hosting providers, DDoS protection services, and registrars, sites nearly always resurface through alternative providers. However, anecdotally, we find that sites struggle to remain online when mainstream monetization channels are severed. We conclude with insights for infrastructure providers and researchers working to stem the spread of misinformation and hateful content.

1 INTRODUCTION

Technical infrastructure providers like Amazon, Cloudflare, and Google have both served and regulated websites that spread misinformation and hateful content. Several influential platforms have extended their service agreements to prohibit such problematic content and, in extreme cases, terminated service to violating sites [15, 33, 56]. For example, in 2017, the *Daily Stormer* lost Distributed Denial of Service (DDoS) protection services from Cloudflare and was subsequently cut off from domain registrar providers GoDaddy and Google, resulting in a website hiatus [11]. Similarly, in 2021, the "far-right alternative to Twitter," Parler, was knocked offline for a month after Apple and Google removed Parler from their app stores and Amazon terminated Parler's hosting services. Yet, despite the growing numbers of infrastructure providers that support misinformation and hate speech websites, there has been little attention paid to identifying who these entities are.

In this paper, we investigate the technical infrastructure that powers misinformation and hate speech websites, including domain registrars, web hosting and email providers, online advertising partners, and DDoS protection providers. We specifically seek to: (1) identify the service providers that misinformation and hate speech websites disproportionately rely on and (2) analyze whether deplatforming such websites affects their long-term availability. To answer these questions, we crawl and analyze the network dependencies of 440 misinformation and hate speech websites—which we refer to in aggregate simply as misinformation websites—from the OpenSources dataset [44]. We crawl each website and collect its DOM, cookies, and network requests, which we then augment with hosting and registrar data. To understand how misinformation websites monetize, we map third-party web dependencies to known advertising providers and payment processors. We then compare misinformation sites to a baseline of mainstream sites.

We show that misinformation sites disproportionately rely on several hosting providers, most prominently Cloudflare, which serves content for 34.3% of misinformation sites compared to 19.6% of mainstream sites. By manually investigating each misinformation website, we find that sites prefer Cloudflare because of its lax acceptable use policies and its free DDoS protection services that help protect against vigilante attacks. Misinformation websites also disproportionately rely on other mainstream providers including GoDaddy, Liquid Web, Sucuri, and Fastly, likely because of their WordPress offerings that allow users to quickly set up and scale sites without much technical expertise. Through a manual analysis of past deplatforming events, we find that when major sites are deplatformed by mainstream hosting and registrar providers, they nearly always find new homes on alternative providers who actively ignore site content, similar to how bullet-proof hosting providers are utilized by malicious actors on the Internet.

Next, we investigate monetization platforms like online advertisement providers and payment processors that enable revenue collection for misinformation sites. A large number of misinformation sites rely on ads—nearly twice the percentage as mainstream sites (62.7% vs 34.9%). Of these misinformation sites, there is significant reliance on mainstream ad networks like Google's DoubleClick (34.4%). Indeed, DoubleClick is used disproportionately by misinformation sites when compared to mainstream sites with ads (effect size of 0.48; 34.4% vs. 14.1% of mainstream sites with ads). In the most severe case, RevContent is used by 22.8% of all misinformation sites but only 0.2% mainstream sites with ads use the provider.

Misinformation sites also disproportionately rely on donations through PayPal and Patreon, as well as direct cryptocurrency donations. While we find little evidence to show that deplatforming by hosting providers is effective at keeping misinformation offline, we note that anecdotally, websites cease producing misinformation content after they are deplatformed from both ad providers and payment processors. In other cases, sites lament the decrease in site revenue after being deplatformed from mainstream ad providers, and as a result, solicit users for direct donations as a means of sustaining site operation.

Conference'17, July 2017, Washington, DC, USA

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

While our investigations began by examining misinformation sites more broadly, misinformation as defined in our dataset is not homogeneous—it is a term encompassing many subcategories for various types of problematic media, including clickbait, conspiracy theories, and particularly, hate speech. Though deplatforming is rare in practice, we find that the misinformation sites that are deplatformed from service providers are deplatformed because of their hateful content (e.g., the *Daily Stormer*, Parler, etc.). We draw upon these case studies to inform us of potentially generalizable solutions and insights that may be effective for misinformation sites more broadly.

We conclude with a discussion of different strategies for preventing the spread of misinformation based on our results. We argue that while deplatforming sites from hosting and registrar infrastructure is likely not an effective solution for combating misinformation, targeting site monetization may be a promising approach. By illuminating what has anecdotally been most effective, we hope to encourage providers—particularly those who have publicized their commitment to fighting misinformation and, more broadly, online abuse—to further explore monetization as a critical channel for curbing the spread of misinformation at scale.

2 RELATED WORK

Our work is inspired by research that highlights the growing complexities of the web. Prior work has studied how websites have grown in complexity [12, 21, 37, 43] and are increasingly relying on centralized network entities and third-party content [37]. Beyond this, several studies have leveraged the nuances of technical infrastructure to better understand and combat traditional computer abuse, including spam and scams [28, 29, 39] and phishing [30, 31]. In particular, Levchenko et al. [39] demonstrate through a series of work that identifying key website infrastructure entities, such as registrar, hosting, and payment providers, helped characterize resource bottlenecks in effective spam intervention.

In the context of misinformation, much work has focused on the classification of websites, primarily through content analysis [38, 48] or social graph features [35, 42, 46, 51]. Studies have leveraged infrastructure properties (e.g., HTTPS configuration or domain expiration) to classify misinformation sites [32], but these studies do not consider web resources broadly as features.

Most closely aligned with our work are several recent studies of the web infrastructure components of misinformation sites. Zeng et al. investigated the ads and ad platforms that power mainstream and misinformation sites, finding that although some advertisers are more prevalent on misinformation sites, both categories share similar fractions of problematic advertising content [57]. Similarly, Agarwal et al. explored the web trackers on hyper-partisan, biased websites. They found that right-leaning, hyper-partisan sites track users more aggressively and rely on many third-party networks (e.g., Doubleclick, Taboola, AdNexus) to function [1].

3 METHODOLOGY

Our study investigates the technical infrastructure that supports misinformation websites, including web hosting, domain registration, DDoS protection, online ads, and payment processing. In this section, we describe the set of misinformation websites we analyze and how we collect data about the providers that support each website.

Misinformation Websites. We analyze a set of misinformation sites that include hate speech sites, which often peddle misinformation; in aggregate, we refer to them as misinformation sites. In this context, we deem misinformation to be non-satirical websites that have potentially misleading content (e.g., "fake news"), determined by the OpenSources project [44]. OpenSources publishes lists of known, vetted, and labeled misleading websites by analyzing sites across several axes: (1) domain name, (2) "About Us" page, (3) article source, (4) writing style, (5) page and image aesthetic, and (6) social media network; these sites predominantly produce content in English, and the set has been used extensively in prior research [10, 32, 50, 57]. While the OpenSources master list contains 826 websites, this list was published in 2017, and because of this, many of these sites are unavailable today. Thus, we removed 191 unreachable sites and 123 parked domains (e.g., those that pointed back to a domain registrar). As our objective is to exclude satirical sites, two independent researchers then manually coded the remaining websites to identify 72 satirical websites, based on the "About pages" of each website in question. Our final misinformation corpus contains 440 websites. The misinformation corpus spans several flavors of unreliability, according to the labels provided by OpenSources. For example, some sites consistently present extreme bias (e.g., breitbart.com), peddle conspiracy theories or bigoted propaganda (e.g., infowars.com, barenakedislam.com), or promote junk science (e.g., naturalnews.com).

Mainstream Website Sample. To construct a baseline of sites to compare misinformation sites with, we consider three candidate sets of sites: (1) 10K random sites from the Alexa Top Million [2], (2) 10K random sites from Certificate Transparency (CT) logs [34], and (3) a list of 579 mainstream news sites [27]. Across the main axes of analysis in this study—hosting and online ad providers—we performed two-sample proportion tests between each of them and our misinformation corpus. In every case, we note that there are statistically significant differences. For instance, for each of the comparisons, Cloudflare was the most disproportionately represented in misinformation websites, with effect size 0.31 in Alexa (Table 1), 0.56 in CT (Table 2), and 0.44 in mainstream news (Table 3).

Because there are significant differences regardless of which set we designate as "mainstream," we choose the 10K random sample from the Alexa Top Million as our baseline, as we believe it best captures the variance in structure and complexity of the sites presented in our misinformation corpus. We explicitly choose not to compare to mainstream news sites, as they are often well-managed and optimized compared to smaller websites that are similar in scale to those in our corpus. Mainstream news sites rely on custom content management systems instead of free-tier systems like WordPress, which is heavily relied on by misinformation websites. We note that 154 (35%) of our misinformation sites appear in the Alexa Top Million; we exclude these from our mainstream corpus.

Data Collection. We crawled 105K pages from August to September 2021. For each page, we allot 10 seconds to navigate to the URL and wait 10 seconds for dynamic content to load. We then collect (1) the page DOM, (2) cookies, and (3) logs of network events. We crawled each website using Crawlium, a crawler

AS	% Misinfo	% Mainstream	p-value	Effect Size
Cloudflare	34.3%	19.6%	1.4e-10	0.31
GoDaddy	6.1%	0.16%	1.9e-07	0.12
Unified Layer	4.7%	0.04%	3.3e-6	0.1
Liquid Web	3.0%	0.1%	0.003	0.05
DigitalOcean	2.7%	0.6%	0.005	0.04

Table 1: ASes Disproportionately Hosting Misinformation Compared to Mainstream Sites—The results of a two-sample proportion test of ASes sorted by effect size. We find that Cloudflare disproportionately represented on misinformation sites.

AS	% Misinfo	% CT	p-value	Effect Size
Cloudflare	34.3%	6.3%	2.5e-33	0.56
GoDaddy	6.1%	1.1%	1.7e-05	0.10
Liquid Web	3.0%	0.6%	0.004	0.05
Sucuri	1.6%	0.1%	0.015	0.03
Fastly	2.0%	0.6%	0.044	0.03

Table 2: ASes Disproportionately Hosting Misinformation Compared to CT Sites—The results of a two-sample proportion test of ASes sorted by effect size. We find that Cloudflare disproportionately represented on misinformation sites.

based on headless Chrome [5]. We visit each website using a fresh browser instance with no cookies. To identify potentially hidden resources or infrastructure on each website, we additionally spider to four first-party links on the same domain.

Resource Analysis. To understand resource dependencies, we construct an inclusion tree for every domain. An inclusion tree is derived from a webpage's DOM and represents the sequence of resource requests made as the site loads its content. We annotate each resource with the origin AS from which it is loaded. We use the AS of the root page to determine each site's web hosting and DDoS protection provider. Narrowing in on the ad provider ecosystem, we determine the entities responsible for image resource loads larger than a 1×1 pixel through domain-entity mappings by WhoTracks.me [55]. A limitation of this approach is that while we restricted our ad detection method to observe images served from ad domains excluding tracking pixels, other images like Facebook's "Like" button are still counted toward Facebook's presence as an ad provider on a site. Finally, we performed a WHOIS lookup on each domain to determine domain registrar and an MX lookup to identify e-mail provider.

Ethical Considerations. We visit each site in our study five times. While this is negligible load for widely-known websites, there are ethical considerations at play as there are with any active scanning. We followed the best practices defined by Durumeric et al. and refer to their work for more detailed discussion of the ethics of active network research [20]. We do not block ads loading because they are an element of our study, but we never click or interact with ads. We argue that we do not significantly impact the misinformation ecosystem along two axes: (1) we do not meaningfully contribute to site traffic in a way that may negatively affect the site itself, and (2) we only negligibly contribute to the ad revenue of misinformation publishers.

4 HOSTING AND DDOS PROTECTION

We first consider the primary hosting provider for each website. We note that because many sites are protected by DDoS protection providers like Cloudflare, in some cases, we can only determine the DDoS protection provider and not the backend hosting provider. In those instances, we classify the site as being hosted by the DDoS provider, since they are still the entities responsible for serving web content to users.

A handful of providers host a disproportionate number of misinformation websites, most notably Cloudflare. Cloudflare provides free CDN and DDoS protection services to sites and is a popular

AS	% Misinfo	% News	p-value	Effect Size
Cloudflare	34.3%	13.0%	1.0e015	0.44
Unified Layer	6.1%	0.5%	6.0e-05	0.09
GoDaddy	3.0%	2.1%	0.002	0.08
OVH SAS	2.7%	0%	4.0e-4	0.05
HIVELOCITY	1.1%	0%	0.02	0.02

Table 3: ASes Disproportionately Hosting Misinformation Compared to Mainstream News Sites—The results of a two-sample proportion test of ASes sorted by effect size. We find that Cloudflare disproportionately represented on misinformation sites.

AS Owner	Sites	% Mis.	AS Owner	Sites	% Mis.
Cloudflare	151	34.3%	Liquid Web	13	3%
Amazon.com	29	6.6%	OVH SAS	12	2.7%
GoDaddy.com	27	6.1%	DigitalOcean	12	2.7%
Google	22	5%	Automatticc	11	2.5%
Unified Layer	21	4.7%	Fastly	9	2%

Table 4: Top Misinformation ASes—We show the top ten ASes responsible for hosting misinformation sites and the portion of misinformation sites for which each is responsible. We find that Cloudflare has the largest market share.

provider across the web, serving 20% of sites in our mainstream sample. However, it also serves the largest fraction of misinformation sites (151 domains, 34.3%) (Table 1). Misinformation sites also disproportionately rely on GoDaddy, Unified Layer, Liquid Web, and Sucuri compared to mainstream sites. To measure this, we conducted a two-sample proportion test, measuring whether the proportion of websites in our misinformation and our mainstream corpus hosted by each hosting provider differed between the two sets. Because we were simultaneously measuring multiple comparisons, we corrected our p-values with Bonferroni corrections $ps < 6.02 \times 10^{-5}$). Given our large sample size, most p-values are statistically significant, so we compute effect size using Cohen's h to better understand the strength of the relationship between these hosting providers and misinformation sites. Our analysis shows that Cloudflare has the largest effect size (0.31, 34.4% of misinformation vs. 19.6% of mainstream sites).

Though not responsible for a disproportionate number of sites, several reputable hosting providers, including Google and Amazon, provide critical infrastructure to dozens of misinformation websites, partially contributing to the global misinformation predicament

Cloudflare Sites	Attack	Cloudflare Migration	
		Date	Post-Attack
barenakedislam.com	2/4/15	8/31/17	✓
drudgereport.com	12/30/16	1/4/17	✓
frontpagemag.com	3/23/15	3/24/15	✓
godlikeproductions.com	4/13/16	8/9/17	✓
naturalnews.com	8/8/17	8/8/17	✓
off-guardian.org	9/26/19	5/6/19	✗
returnofkings.com	9/2/15	10/23/14	✗
russia-insider.com	4/11/18	4/13/18	✓
thegatewaypundit.com	4/15/18	6/12/15	✓
weaselzuppers.us	1/5/15	1/1/14	✗
infostormer.com	12/7/19	8/15/17	✗

Table 5: DDoS Attacks Against Cloudflare Misinformation Sites—Misinformation sites with known DDoS attack history and when they were first observed using Cloudflare hosting in our dataset.

(Table 4). We also find websites protected by well-known CDNs Akamai (e.g., unclesamsmisguidedchildren.com, an extremist site known for its consistent publication of conspiracy theories) and Fastly (e.g., cnsnews.com, an website known for unreliable claims). Across all providers, we find misinformation served from 90 distinct ASes.

4.1 Acceptable Use Policies

In spite of growing concerns regarding misinformation, most hosting providers do not explicitly prohibit hate speech or misinformation. Providers such as GoDaddy, Amazon, Unified Layer, WordPress, and Fastly, do explicitly disavow sites that incite violence, but their terms of service (ToS) and acceptable use policies (AUP) do not extend to hate speech or misinformation [3, 6, 23, 24, 40]. Two hosting providers, OVH and Digital Ocean, specifically prohibit harassing or abusive content, including racially or ethnically offensive content [18, 45]. In contrast, a handful of companies have taken a counter, "content-neutral" approach, notably Cloudflare, whose ToS simply state that it "cannot remove material from the Internet that is hosted by others" [13].

The OpenSources dataset tags each website with additional labels, one of which is whether the website contains hate speech. In our corpus, 30 websites are labeled as hate speech. We find that hateful websites do in fact appear on providers that prohibit the practice. One such website is hosted on OVH (vdare.com), and another by Digital Ocean (actforamerica.org). The most prominent provider among sites specifically serving hateful content is Cloudflare (9 sites, 30%), followed by GoDaddy and Sucuri (3, 10% each).

4.2 Cloudflare DDoS Protection

It is difficult to ascertain exactly why misinformation websites prefer Cloudflare over other hosting providers. Cloudflare has only relatively recently emerged as the primary provider for misinformation and abusive websites (Figure 1). Leveraging historical passive DNS data from Farsight [22], we find that GoDaddy was the most prevalent provider between 2010–2015, hosting up to 48 (11.2%) misinformation sites as recently as 2015. It was not until October 2015 that Cloudflare overtook GoDaddy.

One explanation for Cloudflare's rise is simply that Cloudflare grew in popularity across the Internet. However, we observe that rate of growth for mainstream sites adoption Cloudflare hosting is approximately half that of misinformation websites (Figure 1). In the end, it is likely due to a confluence of reasons. First, it is likely that misinformation websites turn to Cloudflare due to their lax policies. We observe anecdotal evidence from misinformation sites about their reliance on Cloudflare. For example, AmmoLand, a popular guns rights blog, revered Cloudflare not just for its DDoS protection, but also for its self-described "content-neutral" stance [4]:

> *"Cloudflare is also pro-free speech and anti-censorship. Prince is a rarity in Silicon Valley. The SPLC and various left-wing organization have called out Cloudflare to stop providing services to websites that host content that they see as objectionable Cloudflare has responded in a way that I wish more companies would return to this type of pressure from SPLC type groups. They ignored the demands. Prince believes it is imperative for our country that his company remains content-neutral."*

Similarly, an author of Infostormer, who previously wrote for the Daily Stormer, empathized with Cloudflare's CEO concerning difficult decisions of "ban-hammering" sites from their service [33]:

> *"Obviously I don't like what Prince did [banning the Daily Stormer]. I've been highly critical of him over the past few months since he ordered the ban. He justified banning the site because he thought we were "assholes" and happened to be in a bad mood. As a writer for the Daily Stormer I found this comment to be quite offensive. It was also an abandonment of principles. Up until that time, Cloudflare maintained a neutral stance on content. This was the correct position to have. With that said, I can understand that he was put in a tough position. He had to do what he thought was best for the company at that time."*

Second, many sites turn to Cloudflare for its free DDoS protection services because sites regularly come under attack from "vigilante justice" groups [7, 56]. We manually investigate the 151 misinformation websites hosted by Cloudflare and observe that 23 sites have publicly documented experiencing DDoS attacks. Of those, 11 have specific attack dates. Leveraging Farsight's DNS data, we find that 7 (64%) domains transitioned to Cloudflare after an attack, with 4 (37%) transitioning within days of being attacked (Table 5).

In one example, Natural News (naturalnews.com), a prominent anti-vaccination and conspiracy theory site, came under attack on August 8, 2017. At the time, the site relied on Codero and EasyDNS. Then, on the day that news of the DDoS attack on the website was published, Natural News began its transition to Cloudflare. Similarly, FrontPage Magazine (frontpagemag.com), a site known for its far-right, Islamophobic content, experienced a DDoS attack on March 23, 2015. While the site briefly used Cloudflare in May 2013, it quickly switched to using Rackspace Cloud Service's name servers (stabletransit.com). We do not observe changed DNS data until the day after the attack, March 24, 2015, when the site switched to Cloudflare. Both sites have remained on Cloudflare since their respective attacks.

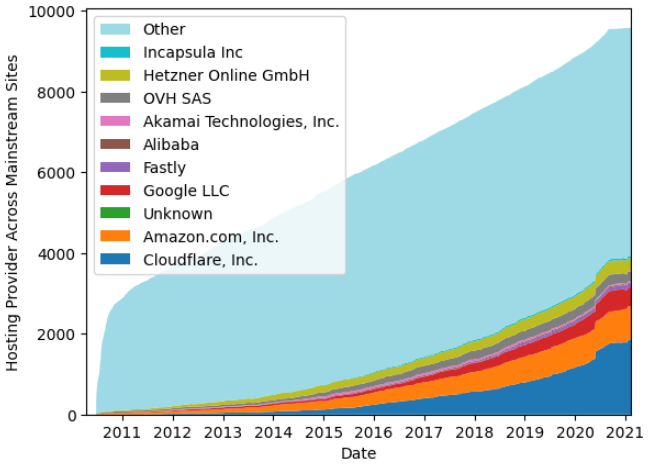

(a) Mainstream Hosting over Time

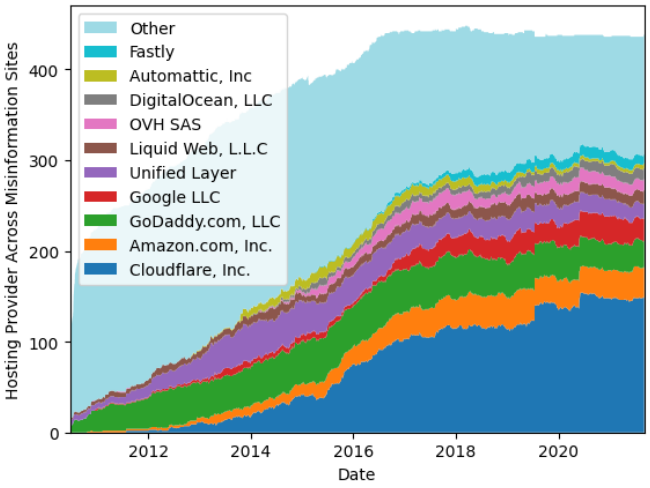

(b) Misinformation Hosting over Time

Figure 1: Longitudinal Hosting Providers—Since the beginning of Farsight's historical DNS data, Cloudflare has seen the most growth of any other hosting provider among both misinformation and mainstream sites. However, misinformation websites have grown far more reliant on Cloudflare, which accounts for 34% of misinformation websites compared to just 17% of mainstream websites.

Registrar	Misinfo	Mainstream	p-value	Effect Size
GoDaddy	42%	24%	3.5e-16	0.39
NameCheap	7.1%	0.8%	1.6e-28	0.35
Epik	3.6%	0.08%	6.3e-39	0.32
eNom	5.7%	1.7%	2.6e-8	0.22
Tucows	4.1%	1.7%	5e-4	0.15
CloudFlare	2.0%	0.6%	3e-4	0.14
NameSilo	1.3%	0.2%	1e-4	0.13
FastDomain	2.6%	1%	0.003	0.12
DNC	1.0%	0.2%	0.003	0.12

Table 6: Registrars Disproportionately Supporting Misinformation—The results of a two-sample proportion test of domain registrars sorted by effect size. GoDaddy is the most prevalent, likely in part due to its free WordPress integration.

Ad Tracker	% Sites	Ad Tracker	% Sites
Facebook	23.4	Outbrain	5.0
DoubleClick	21.6	Taboola	3.9
RevContent	14.3	ShareThis	2.0
Google Syndication	9.1	Connatix	2.0
Google	6.1	Amazon Advertising	1.8

Table 7: Top Advertisers on Misinformation—The distribution of the top 10 advertising trackers found on misinformation sites. Google constitutes three of the 10 advertising domains.

4.3 Hosting and Site Generation Bundles

Many misinformation websites rely on free content management tools. For example, WordPress powers 68% of misinformation websites—nearly twice the percentage of mainstream websites. The reliance on free website generation tools likely explains the prevalence of specific providers. For example, GoDaddy heavily advertises free WordPress integration. About 65% of GoDaddy, 72% of Unified Layer, and 90% of Liquid Web websites use WordPress, highlighting the role that ease of use can play in choosing a hosting provider and supporting misinformation more broadly.

4.4 Domain Registrars

The misinformation sites in our study rely on 47 domain registrars (Table 6). Sites disproportionately rely on GoDaddy (42% misinformation vs. 24% mainstream), NameCheap (7.1% vs. 0.8%), Epik (3.6% vs. 0.08%), eNom (5.7% vs. 1.7%), and Tucows (4.1% vs. 1.7%). We visited the abuse reporting pages of each registrar, and find that while all registrars have abuse reporting mechanisms, only one explicitly prohibits misinformation: Tucows.

In spite of a lack of policy, registrars have made ad-hoc decisions to deplatform hateful, violent, or misleading content in the past. Most notably, the Daily Stormer was deplatformed by a series of registrars, including GoDaddy, Google, and Namecheap, which hindered its ability remain online [49]. Asia Registry, an Australian registrar, booted Gab (an alt-right alternative to Twitter) off of their service in 2017, citing Australian discrimination law [8]. After losing service, Gab switched to Epik, which serves as the registrar for 3.6% of the domains in our misinformation corpus and is disproportionately relied on by misinformation websites. Epik is primarily known for hosting far-right extremist content, famous for previously offering protection services through its company Bitmitigate to 8chan and Parler [9, 26]. Our results highlight lax registrar policies, but also show that many lesser-known registrars (e.g., Epik) are willing to support abusive websites in the name of a free and open Internet.

5 MONETIZATION STRATEGIES

Beyond hosting infrastructure, misinformation sites also rely on online advertising and direct donations to stay online. In this section,

Ad Provider	Misinformation with Ads	Mainstream with Ads	Effect Size
RevContent	22.8%	0.2%	0.91
DoubleClick	34.4%	14.1%	0.48
Outbrain	8.0%	1.5%	0.32
AppNexus	2.2%	0.0%	0.39
Google Syndication	14.5%	6.1%	0.28

Table 8: Advertisers Disproportionately Supporting Misinformation—The top 5 advertising trackers that are found disproportionately often in misinformation over mainstream sites with ads, ordered by effect size. All p-values were were also Bonferroni corrected ($n = 132$) and statistically significant.

we analyze the role monetization plays in supporting misinformation websites.

5.1 Advertising

Online ads continue to be a primary monetization strategy for misinformation websites. Despite journalists calling on ad companies to discontinue serving ads on misinformation sites [52, 53], the majority of misinformation sites use online advertisements at nearly double the rate of mainstream sites (62.7% vs. 34.9%). Many of the ads are served by mainstream providers; for instance, Google's DoubleClick is used on 21.6% of all misinformation sites (Table 7).

Compared to all mainstream sites with ads, misinformation sites with ads disproportionately rely on several ad providers, most notably RevContent and DoubleClick (Table 8). We conducted two-sample proportion tests on the prevalence of all ad providers on misinformation and mainstream websites with ads. All comparisons were corrected for multiple testing using the Bonferroni corrections. To make these comparisons as fair as possible to these providers, we restrict our computation of p-values and effect size to consider only misinformation and mainstream sites with known ad provider dependencies since a higher percentage of misinformation sites have ads.

RevContent. RevContent has the highest effect size (0.91), indicating that it is most disproportionately used by misinformation websites (22.8% of misinformation websites but only 0.2% of mainstream websites with ads). RevContent was previously admonished by mainstream media outlets for serving ads on fake news sites, and even went as far as launching a *Truth in Media Initiative*, which allows users to report misinformation websites. Despite this, the company continues to place ads on known misinformation sites, and the company later defended their inaction, indicating that while fake news intended to deceive was not allowed on the site, satirical content content is not prohibited [52]. We note that we removed satirical sites from our misinformation corpus; 15.9% of sites using RevContent in our study are labeled as junk science sites and 34.9% as conspiracy theory sites.

DoubleClick. Googles's DoubleClick has the second highest effect size (0.48), present in 95 (34.4%) misinformation sites with ads. In response to rising concern over misinformation amidst the 2016 U.S. election, Google released a statement that it would "restrict ad serving on pages that misrepresent, misstate, or conceal information about the publisher, the publisher's content, or the primary purpose of the web property" [41]. According to our dataset's site labels, however, of the sites serviced by DoubleClick, 27 (28.4%) are conspiracy sites; 17 (28.4%) are fake news; 9 (9.5%) are junk science.

There is anecdotal evidence that removing ad revenue is effective at curbing the spread of misinformation. In one such instance, Google AdSense deplatformed American Free Press in 2017 for serving anti-Semitic content. The site remains blocked by Google Ads. Today, most ads found on American Free Press are embedded directly into the page as first-party content. We also detect the header bidding library, Prebid.js, on American Free Press, allowing the site to directly offer bid slots to brands. American Free Press has experienced a different fate from that of ZeroHedge, which was deplatformed by Google in June 2019. ZeroHedge's ad monetization was reinstated by Google just one month later after its takedown of problematic comments [25]. All News Pipeline, a conspiracy theory site, laments the decline of ad revenue for itself and other "independent" media sites [19]:

> *"With digital media revenue spiraling downward, especially hitting those in Independent Media, it has become apparent that traditional advertising simply isn't going to fully cover the costs and expenses for many smaller independent websites."*

This hints that ad providers may be able to effectively reduce misinformation driven revenue and site operation. However, we note that sites rely on an average of seven ad providers, underscoring the need for coordinated efforts amongst providers. Unfortunately, this does not appear to be happening in practice. Despite RevContent and Google previously claiming to be curbing misinformation on their platforms [41, 52, 53], our results indicate that their efforts are not effective, and that these organizations still financially support the spread of misinformation. Broadly, we find minimal evidence of ad providers blocking misinformation sites.

5.2 Donations

Misinformation websites often also rely on donations to sustain their operations. Donation strategies range from using third-party intermediaries like PayPal to solicit donations to directly accepting cryptocurrencies like Bitcoin. In our corpus, 43 (9.78%) misinformation sites rely on resources from PayPal compared to only 67 (0.01%) mainstream websites. A two-sample proportion test indicates that this difference in proportions is statistically significant ($p < 0.005$, $h = 0.47$): misinformation sites disproportionately rely on PayPal compared to mainstream sites.

To understand why PayPal is disproportionately represented on misinformation sites, we manually investigated the 43 misinformation sites that loaded resources from PayPal domains. For each of these sites, we examined web pages and banners soliciting donations. We find that 93% (40) of the sites that rely on PayPal use it for donation services, but two links were inactive. The remainder (7%) utilized PayPal for subscriptions or storefronts. The misinformation sites in our investigation also solicit Bitcoin donations (14%), Patreon donations (9.3%), and Salsa Labs donations (4.7%).

PayPal has previously blocked payment on sites hosting hateful and non-inclusive content. One site author that was deplatformed by PayPal is Roosh Valizadeh (Roosh V) known for his support of men's rights and the alt-right. One of his sites, returnofkings.com,

AS	ASN	Misinfo		Mainstream	
		#	%	#	%
Google	13949	416	94.5%	7973	79.7%
Amazon.com	14618	324	73.6%	3771	37.7%
Cloudflare	13335	308	70%	3804	38%
Fastly	54113	283	64.3%	2396	24%
Akamai	393234	239	54.3%	2293	22.9%
Facebook	32934	228	51.8%	3262	32.6%
AppNexus	36805	199	45.2%	1129	11.3%
Highwinds	11588	198	45%	2138	21.4%
MCI	12199	182	41.4%	1096	11%
Automattic	2635	175	39.8%	858	8.6%

Table 9: Top Misinformation Third-Party Resource ASes—The top 10 autonomous systems responsible for third-party resource loads across misinformation sites.

is present in our set of misinformation sites. Return of Kings (ROK) announced a hiatus in 2018, identifying deplatforming of monetization strategies (e.g., PayPal and ads) as a successful tactic in removing misinformation online:

> *"The first factor for this hiatus is that site revenues are too low. We've been banned from Paypal and countless ad partners, which forced me to lay off the site editor last year and also lower payments to regular contributors. This started a negative spiral of declining content quality, site traffic, and revenues. Even the beloved comments section, which many see as the highlight of ROK, was badly hit when Disqus banned us. Currently, ROK receives half the traffic of its peak and less than one-fifth of the income"* [54].

The Daily Stormer also faced challenges from restricted revenue streams, but remains operational. As a result, it has been forced to rely solely on donations:

> *"We are not allowed to use any form of advertisement. We cannot use PayPal. We cannot even use credit card processors. We had a P.O. box, and even that was shut down. The only way we can receive money is through crypto currency"* [16].

Our data indicates a variety of different revenue streams supporting the production of misleading content online, but hints that coordinated deplatforming by both ad providers and payment processors may be an effective way of disincentivizing the continued upkeep of online misinformation.

6 OTHER TECHNICAL DEPENDENCIES

Although hosting platforms and monetization sources are the primary dependencies for misinformation sites, sites often rely on a myriad of other technical dependencies like third-party web resources and e-mail providers. In this section, we highlight these other dependencies.

Misinformation websites, which can range from complex news pages to small blogs, load a median of 215 resources, of which 77 (36%) are first-party and 138 (64%) are third-party. Compared to mainstream sites, which rely on a median of 36% third-party resources, misinformation sites more heavily rely on third-party entities. Misinformation sites load the same top third-party resources as mainstream sites (e.g., popular analytics, tracking, and advertising resources). In some cases, misinformation sites do have statistically different proportions: for example, 61% of misinformation websites rely on DoubleClick whereas only 35% of mainstream websites do. However, most differences are marginal. The third-party resources that misinformation sites rely on come from a variety of providers; however, a small handful of providers. Unsurprisingly, misinformation sites depend on resources from major players including Google (95%), Amazon (74%), Cloudflare (70%), Fastly (64.3%), and Akamai (54.3%) (Table 9). Beyond previously discussed services (e.g., Google ads), large providers also support website in other manners. For example, Google also provides fonts (83% of misinformation websites) and custom search integration (69% of misinformation websites).

Many misinformation sites are also configured to accept inbound email. We find no statistically significant differences in e-mail providers. Misinformation sites most commonly depend on Microsoft Outlook and Gmail, which serve 32% and 17% of misinformation websites, respectively. Similarly to analytics providers, we see evidence of inconsistent policies within companies. Despite being deplatformed by Google News, westernjounalism.com still uses Gmail. We encourage organizations that make deplatforming decisions to consider all products that may be used to support misinformation operations.

7 DISCUSSION

We find limited evidence of infrastructure providers deplatforming misinformation sites. Among the misinformation sites that were deplatformed, platform providers often cited the violent and hateful nature of the content—rather than its credibility—as the bases for such decisions. Our results suggest that deplatforming hate speech websites from hosting services only has short-term impact on their availability; often, echoing the relationships between spam sites and bulletproof hosting providers examined in prior work, these sites eventually find alternate providers and return online.

While deplatforming sites from hosting providers may not be the most promising avenue, a large number of misinformation and hate speech sites depend on ad providers, including mainstream companies like Google DoubleClick. Among smaller hate speech sites, we find anecdotal evidence that deplatforming sites from monetization channels may have long-term success in stemming the production of new problematic content. It still remains to be seen if these strategies can also be applied to other kinds of misinformation more broadly.

7.1 Hosting and Domain Registration

Though some mainstream hosting and domain registration providers have policies condemning hate and violence, policies against misinformation sites broadly are limited. Furthermore, even when providers have these policies, enforcement is not comprehensive. In the few cases where sites are deplatformed, there are many alternative providers available to site operators. Similar to how bulletproof hosting providers allow customers to host illegal content, send

spam, and launch DDoS attacks [36], niche registrars and hosting companies are willing to serve misinformation and abusive content.

Broadly, we find that deplatforming misinformation from hosting providers does not prevent them from remaining online in the long term. For example, while the Daily Stormer faced issues staying online for over two years after a series of deplatforming instances with hosting and domain registrar services, it was eventually able to find stable hosting with VanwaTech and registrar services with Russian provider R01. For similar reasons, Parler went offline for several weeks, but eventually returned online with Beelastic as its hosting provider. However, we note that alternative providers tend to be more expensive, and because lesser known misinformation sites have not been deplatformed in practice, it remains unclear whether small sites would be able to afford alternatives. Furthermore, while mainstream hosting providers are not in the position to best stem online misinformation, they should reconsider their role in actively supporting such content.

7.2 DDoS Protection

While DDoS protection is not a required component of infrastructure for many mainstream websites, there is a long history of particularly offensive or hateful misinformation sites coming under attack, and empirically DDoS protection is a particularly useful service for these sites. Misinformation websites disproportionately rely on Cloudflare, a provider that offers *free* DDoS protection and has neglected to address abusive content in all but the most egregious cases. We observe steady growth in Cloudflare's popularity across misinformation sites since 2010; today, Cloudflare is the primary provider for misinformation sites.

This may be due to the absence of free alternative DDoS protection. Misinformation websites have only a few alternatives, many of which are expensive: Bitmitigate, which serviced the Daily Stormer after it was removed from Cloudflare, costs $159 a month for enterprise-level protection, and DDoS-Guard, which is leveraged by several far-right websites in our dataset, costs up to $1,000 a month. While both Cloudflare and DDoS-Guard offer a free tier of protection, DDoS-Guard's free tier only offers protection for attacks with up to 1.5 Tbps compared to Cloudflare's 67 Tbps capacity [14, 17]. It remains unclear whether smaller DDoS protection providers—especially at analogous free tiers of service—can withstand significant attacks.

7.3 Monetization

Most misinformation sites depend on online ads, and these sites often rely on mainstream ad providers like Google for these services. A disproportionate number of misinformation and hate sites rely on online advertising from companies that have already pledged to combat misinformation like Google and RevContent. Extrapolating from anecdotal accounts from smaller hate speech sites, blocking monetization channels for sites may be the most promising avenue for curbing the spread of misinformation. Anecdotally, misinformation websites report significant decreases in revenue, and in some cases stop publishing new content entirely after losing access to advertising and donation platforms. For example, the website *Return of Kings* was first deplatformed by PayPal, and eventually shut off by almost all advertising partners. While ads from MGID, a native advertising company, are still displayed on the site, the overall decrease in revenue forced the site to announce a hiatus, and no new content has been posted since October 2018. Aligned with prior research in online abuse that suggests that increasing costs reduces harm [47], we suggest that monetization platforms—both ad providers and payment processors—consider the role they play in supporting online misinformation and how they may be in the best position to curb its spread.

7.4 Ethics of Deplatforming

Our paper focuses on understanding the providers that directly or indirectly support misinformation websites and whether deplatforming helps curb the spread of misinformation. It remains an open question whether companies *should* deplatform all kinds of misinformation sites, and if they do, how they should choose which sites to deplatform. While a few providers have policies that prohibit misinformation, many do not, which may inadvertently enable misinformation websites to thrive on their platforms. We encourage providers to actively consider writing concrete policies around abusive content and misinformation. We also note that several of the largest ad providers have publicly announced their intent to fight online abusive content and misinformation; however, according to our data, they have failed to take meaningful action against such sites. For instance, over 20% of *all* misinformation sites rely on Google for ads. These mainstream ad providers are not only supporting misinformation sites by providing them ad revenue, but also profiting from maintaining relationships with these publishers. We encourage providers to reconsider how they are enforcing their policies.

8 CONCLUSION

In this paper, we analyzed the infrastructure that powers misinformation websites. We showed that several providers are disproportionately responsible for hosting online misinformation, most prominently Cloudflare, which hosts a third of the misinformation sites in our study. While many providers prohibit hateful content, they rarely have clauses in their terms of service to forbid general misinformation on their platforms, and even when hosting providers do have policies, enforcement is rare and seemingly ineffective. When misinformation websites are deplatformed by hosting providers and registrars, they find other willing providers to serve their content. However, we do find that misinformation sites disproportionately rely on monetization platforms like ad networks and donation platforms, and that anecdotally, sites appear to struggle when their monetization channels are removed. We hope our results will inform infrastructure platforms and researchers of more effective strategies to reduce the spread of online misinformation.

REFERENCES

[1] Pushkal Agarwal, Sagar Joglekar, Panagiotis Papadopoulos, Nishanth Sastry, and Nicolas Kourtellis. 2020. Stop tracking me Bro! Differential Tracking of User Demographics on Hyper-Partisan Websites. In *The World Wide Web Conference*.

[2] Alexa Internet, Inc. 2020. Top 1,000,000 Sites. http://s3.amazonaws.com/alexa-static/top-1m.csv.zip. (2020). 2020-10-08.

[3] Amazon Web Services, Inc. 2016. AWS Acceptable Use Policy. https://aws.amazon.com/aup/. (2016). Accessed: 2021-01-15.

[4] AmmoLand. 2020. Gun Owner Privacy. https://www.ammoland.com/tags/gun-owner-privacy/feed/. (2020).

[5] Sajjad Arshad. 2020. Crawlium. https://github.com/sajjadium/Crawlium. (2020). Accessed: 2020-10-08.

[6] Automattic. 2021. Terms of Service. https://wordpress.com/tos/. (2021). Accessed: 2021-01-15.

[7] BBC. 2017. Daily Stormer: Cloudflare drops neo-Nazi site. (2017). https://www.bbc.com/news/technology-40960053

[8] Ali Breland. 2017. Alt-right Twitter rival may lose its web domain. https://thehill.com/policy/technology/351169-alt-right-twitter-rival-might-lose-its-domain. (2017).

[9] Jon Brodkin. 2019. Dumped by Cloudflare, 8chan gets back online—then gets kicked off again. *Ars Technica* (2019). https://arstechnica.com/tech-policy/2019/08/8chan-briefly-got-back-online-with-same-cdn-used-by-neo-nazi-daily-stormer/

[10] Ceren Budak. 2019. What Happened? The Spread of Fake News Publisher Content During the 2016 U.S. Presidential Election. In *The World Wide Web Conference (The World Wide Web Conference)*. Association for Computing Machinery, New York, NY, USA, 139–150. https://doi.org/10.1145/3308558.3313721

[11] Sean Burch. 2017. https://www.thewrap.com/russian-internet-boot-daily-stormer/. (2017). Accessed: 2021-01-15.

[12] Michael Butkiewicz, Harsha V Madhyastha, and Vyas Sekar. 2011. Understanding website complexity: measurements, metrics, and implications. In *ACM SIGCOMM: Internet Measurement Conference*.

[13] Cloudflare. 2020. Website and Online Services Terms of Use. https://www.cloudflare.com/website-terms/. (2020).

[14] Cloudflare. 2021. https://www.cloudflare.com/plans/. (2021). Accessed: 2021-01-15.

[15] Kate Cox. 2021. Parler goes dark, sues Amazon to demand immediate reinstatement. *Ars Technica* (2021). https://arstechnica.com/tech-policy/2021/01/parler-goes-dark-sues-amazon-to-demand-immediate-reinstatement/

[16] DailyStormer. 2021. Support the Daily Stormer! https://dailystormer.su/contributions/. (2021). Accessed: 2021-01-15.

[17] DDoS-Guard. 2021. https://ddos-guard.net/en/store/web. (2021). Accessed: 2021-01-15.

[18] DigitalOcean. 2020. Acceptable Use Policy. https://www.digitalocean.com/legal/acceptable-use-policy/. (2020).

[19] Susan Duclos. 2018. Independent Media Labeled 'Dangerous' & Disqus Comment Notifications Listed In Gmail As Suspicious Or Spam As Big Tech Finds News Ways To Attack Alternative Media Websites. https://allnewspipeline.com/Tech_Attacks_On_IM_Ramps_Up.php. (2018). Accessed: 2021-01-15.

[20] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. 2013. ZMap: Fast Internet-Wide Scanning and its Security Applications. In *USENIX Security Symposium*.

[21] Steven Englehardt and Arvind Narayanan. 2016. Online tracking: A 1-million-site measurement and analysis. In *ACM SIGSAC: Conference on Computer and Communications Security)*.

[22] Farsight Security. 2021. Introducing DNSDB 2.0. https://www.farsightsecurity.com/solutions/dnsdb. (2021). 2021-01-11.

[23] Fastly, Inc. 2021. ACCEPTABLE USE POLICY; REPORTING A VIOLATION; AND DMCA SAFE HARBOR. https://www.fastly.com/acceptable-use/. (2021). Accessed: 2021-01-15.

[24] GoDaddy. 2020. GoDaddy Legal Agreements and Policies. https://www.godaddy.com/legal/agreements. (2020).

[25] Megan Graham. 2020. Google says Zero Hedge can run Google ads again after removing 'derogatory' comments. https://www.cnbc.com/2020/07/14/google-reinstates-zero-hedge-ad-monetization.html. (2020).

[26] Rachel E. Greenspan. 2021. Parler moves to Epik, a domain registrar known for hosting far-right extremist content. https://www.businessinsider.com/parler-moves-to-epik-domain-known-for-hosting-far-right-2021-1. (2021).

[27] Hans W.A. Hanley, Deepak Kumar, and Zakir Durumeric. 2022. No Calm in the Storm: Investigating QAnon Website Relationships. In *AAAI Conference on Web and Social Media*.

[28] Shuang Hao, Alex Kantchelian, Brad Miller, Vern Paxson, and Nick Feamster. 2016. PREDATOR: Proactive Recognition and Elimination of Domain Abuse at Time-Of-Registration. In *ACM SIGSAC: Conference on Computer and Communications Security)*.

[29] Shuang Hao, Nadeem Ahmed Syed, Nick Feamster, Alexander G Gray, and Sven Krasser. 2009. Detecting Spammers with SNARE: Spatio-temporal Network-level Automatic Reputation Engine.. In *USENIX Security Symposium*.

[30] Grant Ho, Asaf Cidon, Lior Gavish, Marco Schweighauser, Vern Paxson, Stefan Savage, Geoffrey M Voelker, and David Wagner. 2019. Detecting and characterizing lateral phishing at scale. In *USENIX Security Symposium*.

[31] Grant Ho, Aashish Sharma, Mobin Javed, Vern Paxson, and David Wagner. 2017. Detecting credential spearphishing in enterprise settings. In *USENIX Security Symposium*.

[32] Austin Hounsel, Jordan Holland, Ben Kaiser, Kevin Borgolte, Nick Feamster, and Jonathan Mayer. 2020. Identifying Disinformation Websites Using Infrastructure Features. In *USENIX Workshop on Free and Open Communications on the Internet*.

[33] Infostormer. 2019. Infostormer Has Been Under a Sustained DDos Attack This Week. https://infostormer.com/infostormer-has-been-under-a-sustained-ddos-attack-this-week/. (2019).

[34] Internet Engineering Task Force. 2013. https://datatracker.ietf.org/doc/html/rfc6962. (2013). Accessed: 2021-01-15.

[35] Zhiwei Jin, Juan Cao, Yu-Gang Jiang, and Yongdong Zhang. 2014. News credibility evaluation on microblog with a hierarchical propagation model. In *IEEE International Conference on Data Mining*.

[36] Maria Konte, Roberto Perdisci, and Nick Feamster. 2015. Aswatch: An as reputation system to expose bulletproof hosting ASes. In *ACM SIGCOMM: Conference on Special Interest Group on Data Communication*.

[37] Deepak Kumar, Zane Ma, Zakir Durumeric, Ariana Mirian, Joshua Mason, J Alex Halderman, and Michael Bailey. 2017. Security challenges in an increasingly tangled web. In *The World Wide Web Conference*.

[38] Srijan Kumar, Robert West, and Jure Leskovec. 2016. Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes. In *The World Wide Web Conference*.

[39] Kirill Levchenko, Andreas Pitsillidis, ha Chachra, Brandon Enright, Márk Félegyházi, Tristan Halvorson, Chris Kanich, Christian Kreibich, He Liu, Damon McCoy, Nicholas Weaver, Vern Paxson, Geoffrey M. Voelker, and Stefan Savage. 2011. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *IEEE Symposium on Security and Privacy*.

[40] Liquid Web. 2021. Liquid Web Acceptable Use Policy ("AUP"). https://www.liquidweb.com/about-us/policies/acceptable-use-policy/. (2021). Accessed: 2021-01-15.

[41] Julia Love and Kristina Cooke. 2016. Google, Facebook move to restrict ads on fake news sites. *Reuters* (2016). https://www.reuters.com/article/us-alphabet-advertising/google-facebook-move-to-restrict-ads-on-fake-news-sites-idUSKBN1392MM

[42] Van-Hoang Nguyen, Kazunari Sugiyama, Preslav Nakov, and Min-Yen Kan. 2020. FANG: Leveraging social context for fake news detection using graph representation. In *ACM SIGIR: Conference on Information & Knowledge Management*.

[43] Nick Nikiforakis, Luca Invernizzi, Alexandros Kapravelos, Steven Van Acker, Wouter Joosen, Christopher Kruegel, Frank Piessens, and Giovanni Vigna. 2012. You are what you include: large-scale evaluation of remote javascript inclusions. In *ACM SIGSAC: Conference on Computer and Communications Security)*.

[44] OpenSources. 2017. OpenSources. https://github.com/bigmclargehuge/opensources. (2017). Accessed: 2020-10-08.

[45] OVHcloud. 2020. Terms of Service. https://us.ovhcloud.com/legal/terms-of-service. (2020). Accessed 2021-01-15.

[46] Kashyap Popat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2017. Where the truth lies: Explaining the credibility of emerging claims on the web and social media. In *The World Wide Web Conference*.

[47] Anirudh Ramachandran and Nick Feamster. 2006. Understanding the Network-Level Behavior of Spammers. In *ACM SIGCOMM: Conference on Applications, technologies, architectures, and protocols for computer communications*.

[48] Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Conference on Empirical Methods in Natural Language Processing*.

[49] Adi Robertson and Andrew Liptak. 2017. https://www.theverge.com/2017/8/20/16170370/namecheap-host-take-down-neo-nazi-hate-site-daily-stormer. (2017). Accessed: 2021-01-15.

[50] Karishma Sharma, Feng Qian, He Jiang, Natali Ruchansky, Ming Zhang, and Yan Liu. 2019. Combating Fake News: A Survey on Identification and Mitigation Techniques. *ACM Transactions on Intelligent Systems and Technology* (2019).

[51] Kai Shu, Suhang Wang, and Huan Liu. 2019. Beyond news contents: The role of social context for fake news detection. In *ACM International Conference on Web Search and Data Mining*.

[52] Craig Silverman. 2017. An Ad Network That Helps Fake News Sites Earn Money Is Now Asking Users To Report Fake News. *BuzzFeed News* (2017). https://www.buzzfeednews.com/article/craigsilverman/an-ad-network-that-works-with-fake-news-sites-just-launched

[53] Craig Silverman, Jeremy Singer-Vine, and Lam Thuy Vo. 2017. In Spite Of The Crackdown, Fake News Publishers Are Still Earning Money From Major Ad Networks. *BuzzFeed News* (2017). https://www.buzzfeednews.com/article/craigsilverman/fake-news-real-ads

[54] Roosh Valizadeh. 2018. RETURN OF KINGS IS GOING ON HIATUS. https://www.returnofkings.com/195790/return-of-kings-is-going-on-hiatus. (2018). Accessed: 2021-01-15.

[55] WhoTracks.Me. 2021. https://whotracks.me. (2021).

[56] Julia Carrie Wong. 2019. 8chan: the far-right website linked to the rise in hate crimes. *The Guardian* (2019). https://www.theguardian.com/technology/2019/aug/04/mass-shootings-el-paso-texas-dayton-ohio-8chan-far-right-website

[57] Eric Zeng, Tadayoshi Kohno, and Franziska Roesner. 2020. Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites. In *Workshop on Technology and Consumer Protection*.