Proceedings of the Sixteenth International AAAI Conference on Web and Social Media (ICWSM 2022)

On the Infrastructure Providers That Support Misinformation Websites

Catherine Han, Deepak Kumar, Zakir Durumeric
Stanford University
cathan@stanford.edu, kumarde@stanford.edu, zakird@stanford.edu

Abstract

In this paper, we analyze the service providers that power 440 misinformation and hate sites, including hosting platforms, domain registrars, CDN providers, DDoS protection companies, advertising networks, donation processors, and e-mail providers. We find that several providers are disproportionately responsible for serving misinformation websites, most prominently Cloudflare. We further show that misinformation sites 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 hate content.

Introduction

Technical infrastructure providers like Amazon, Cloudflare, and Google have both served and regulated websites that spread misinformation and hate content. Several influential platforms have extended their service agreements to prohibit such problematic content and, in extreme cases, terminated service to violating sites (Cox 2021; Wong 2019; Infostormer 2019). For example, in 2017, neo-Nazi and white supremacist site 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 (Burch 2017). 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 provide technical services for 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 CDN 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 because nearly all hate speech websites also spread misinformation, we refer to in aggregate simply as misinformation websites—from the OpenSources dataset (OpenSources 2017). 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 providers to serve web content, 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 anecdotally 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. We analyze past deplatforming events and 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. The majority of misinformation sites rely on ads—nearly twice the percentage as mainstream sites (62.7% vs. 34.9%). These misinformation and hate sites rely significantly on mainstream ad networks like Google DoubleClick (34.4%). Indeed, DoubleClick is disproportionately used by more than a third of all misinformation sites in our study when compared to

mainstream sites that use ads (effect size of 0.48; 34.4% vs. 14.1% of mainstream sites with ads). In the most severe case, RevContent ads are on 22.8% of all misinformation sites but only 0.2% of mainstream sites with ads.

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.

While our investigations began by examining general misinformation sites, 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 hate 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.

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 (Butkiewicz, Madhyastha, and Sekar 2011; Nikiforakis et al. 2012; Englehardt and Narayanan 2016; Kumar et al. 2017) and are increasingly relying on centralized network entities and third-party content (Kumar et al. 2017). Beyond this, several studies have leveraged the nuances of technical infrastructure to better understand and combat traditional computer abuse, including spam, scams (Hao et al. 2009, 2016; Levchenko et al. 2011), and phishing (Ho et al. 2017, 2019). In particular, Levchenko et al. (Levchenko et al. 2011) 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 (Rashkin et al. 2017; Kumar, West, and Leskovec 2016) or social graph features (Nguyen et al. 2020; Jin et al. 2014; Popat et al. 2017; Shu, Wang, and Liu 2019). Studies have leveraged infrastructure properties (e.g., HTTPS configuration or domain expiration) to classify misinformation sites (Hounsel et al. 2020), but these studies do not consider web resources broadly as features. We incorporate some of these infrastructure properties into our analyses.

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 (Zeng, Kohno, and Roesner 2020). 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 services (e.g., Doubleclick, Taboola, AdNexus) to function (Agarwal et al. 2020).

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. 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 (OpenSources 2017). 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 (Zeng, Kohno, and Roesner 2020; Hounsel et al. 2020; Budak 2019; Sharma et al. 2019). 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). We analyze this set of misinformation sites that include hate speech sites, which often peddle misinformation; in aggregate, we refer to them as misinformation sites. 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 OpenSources' labels. 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 (Alexa Internet, Inc. 2020), (2) 10K random sites from Certificate Transparency (CT) logs (Internet Engineering Task Force 2013), and (3) a list of 579 mainstream news sites (Hanley, Kumar, and Durumeric 2022). The Alexa Top Million is a list of sites curated by Amazon, which ranks them by their one month Alexa traffic rank. Because this sample may be biased by being among the top million most popular sites known to Amazon, we also include the CT sample. CT logs seek to record all certificates publicly issued by certificate authorities; therefore this sample does not suffer from the same notion of popularity that the Alexa Top Million sites do. We additionally considered the mainstream news sites in order to confirm that the differences observed were not due to the type of sites we were measuring (e.g., news) in our mainstream sample. 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 example, for each of the comparisons, Cloudflare was the most disproportionately represented in misinformation websites, with effect size 0.31 in Alexa, 0.56 in CT, and 0.44 in mainstream news (Table 1).

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. Prior work has found that news and sports sites are the most complex in site composition, loading more resources than other kinds of sites due to their reliance on media-based content and their popularity (Kumar et al. 2017). Thus, these differences are exacerbated when we compare these mainstream news sites to the websites in our misinformation corpus. Further, 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. To identify potentially hidden resources or infrastructure, we spider to four first-party links for each of the websites. In total, 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 Document Object Model (DOM), (2) cookies, and (3) logs of network events. The DOM of a page represents its layout as a tree, which we can use to understand the relationship between every element and resource on the page. We crawled each website using a modified version of Crawlium, a crawler based on headless Chrome (Arshad 2020). We visit each website using a fresh browser instance with no cookies.

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 autonomous system (AS) from which it is loaded. Determining the AS helps to identify the entity or organization serving content for a web resource. We use the AS of the root page to determine each site's web hosting or Content Distribution Network (CDN) and DDoS protection provider. The web hosting provider is responsible for storing and serving the contents of the web page to allow for better availability and performance. Because of the complexity of the web, however, sometimes CDN or DDoS protection providers obscure visibility into knowing the hosting 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 (WhoTracks.Me 2021). 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 its 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 (Durumeric, Wustrow, and Halderman 2013). 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.

Hosting and DDoS Protection

We first consider the provider that serves content for each website. We note that because many sites are protected by CDN and DDoS protection providers like Cloudflare, in some cases, we can only determine the forward-facing provider and not the backend hosting provider. For instance, while Cloudflare does not offer servers for publishers to store the content of their sites, once Cloudflare fetches this content, it is served directly to users from Cloudflare infrastructure. Further, the services they offer are critical in serving their content securely and quickly on the web. In addition to DDoS mitigation and CDN services, Cloudflare offers domain name service (DNS) and even domain registrar services——services that web hosts are often responsible for. We refer to these kinds of network service providers generally as hosting providers, since they are among the entities responsible for serving web content to users.

A handful of providers serve content for a disproportionate number of misinformation websites, most notably Cloudflare. Cloudflare is a popular provider across the web,

AS	Misinfo	Mainstream			CT			News		
	%	%	p-val	Eff. Size	%	p-val	Eff. Size	%	p-val	Eff. Size
Cloudflare	34.3	19.6	1.4e-10	0.31	6.3	2.5e-33	0.56	13.0	1.0e-15	0.44
GoDaddy	6.1	0.2	1.9e-7	0.12	1.1	1.7e-5	0.10	2.1	0.002	0.08
Unified Layer	4.7	0.0	3.3e-6	0.10	6.3	0.01	0.03	0.5	6.0e-5	0.09
Liquid Web	3.0	0.1	0.003	0.05	0.6	0.004	0.05	2.0	0.29	0.02
DigitalOcean	2.7	0.6	0.005	0.04	1.4	0.11	0.02	1.7	0.29	0.00
Sucuri	1.6	0.3	0.03	0.03	0.1	0.015	0.03	0.1	0.31	0.01
Fastly	2.0	3.2	0.10	0.02	0.6	0.044	0.03	16.8	3.9e-18	0.00
OVH SAS	2.7	1.2	0.05	0.03	2.6	0.9	0.00	0.0	4.5e-4	0.05
Hivelocity	1.1	0.1	0.04	0.02	0.2	0.07	0.02	0.0	0.02	0.02

Table 1: ASes Disproportionately Serving Misinformation Compared to Other Sites—The results of a two-sample proportion test of ASes between misinformation and mainstream, CT, and mainstream news sites, from left to right. We find that Cloudflare is the most disproportionately represented AS on misinformation sites across all comparisons.

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%	Automattic	11	2.5%
Unified Layer	21	4.7%	Fastly	9	2%

Table 2: Top Misinformation ASes—We show the top ten ASes responsible for serving content for misinformation sites and the portion of misinformation sites for which each is responsible. We find that Cloudflare has the largest market share.

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 served by each 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 providers and misinformation sites. The magnitude of the effect sizes give interpretations for small, medium, and large differences, and these thresholds are often 0.2, 0.5, and 0.8, respectively. Our analysis shows that Cloudflare has the largest effect size (0.31, 34.4% of misinformation vs. 19.6% of mainstream sites), indicating it has the largest difference in proportions.

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 (Table 2). 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.

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 (GoDaddy 2020; Amazon Web Services, Inc. 2016; Automattic 2021; Fastly, Inc. 2021; Liquid Web 2021). Two hosting providers, OVH and Digital Ocean, specifically prohibit harassing or abusive content, including racially or ethnically offensive content (OVHcloud 2020; DigitalOcean 2020). 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" (Cloudflare 2020).

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 hate 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 hate content is Cloudflare (9 sites, 30%), followed by GoDaddy and Sucuri (3, 10% each).

Cloudflare

It is difficult to ascertain exactly why misinformation websites prefer to serve content through Cloudflare. Cloudflare has only relatively recently emerged as the primary provider for misinformation and abusive websites (Figure 1). Leveraging historical passive DNS data from Farsight (Farsight Security 2021), we find that GoDaddy was the most prevalent provider between 2010–2015, hosting up to 48 (11.2%)

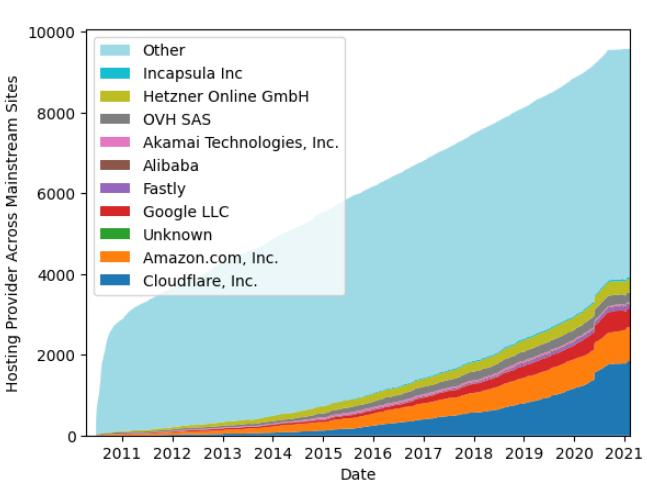

(a) Mainstream Hosting over Time

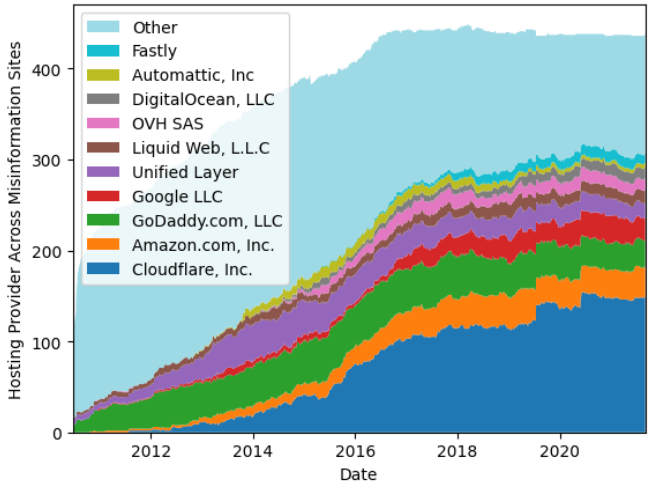

(b) Misinformation Hosting over Time

Figure 1: Longitudinal Hosting Providers—Since the beginning of Farsight's historical DNS data, Cloudflare has seen the most growth of any other provider serving content for both misinformation and mainstream sites. However, misinformation websites have grown far more reliant on Cloudflare, which accounts for 34.3% of misinformation websites compared to just 19.6% of mainstream websites.

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 3: DDoS Attacks Against Cloudflare Misinformation Sites—Misinformation sites with known DDoS attack history and when they were first observed using Cloudflare.

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 adopting 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 (AmmoLand 2020):

> *"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 Matthew Prince concerning difficult decisions of "banhammering" sites from their service (Infostormer 2019):

> *"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 controversial sites regularly come under attack from "vigilante justice" groups (BBC 2017; Wong 2019). 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 known 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 3).

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 4: 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.

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.

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.

Domain Registrars

The misinformation sites in our study rely on 47 domain registrars (Table 4, Table 5). 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 (Robertson and Liptak 2017). Asia Registry, an Australian registrar, booted Gab (an alt-right alternative to Twitter) off of their service in 2017, citing Australian discrimination law (Breland 2017). 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 (Brodkin 2019; Greenspan 2021). 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.

AS Owner	Sites	% Mis.	AS Owner	Sites	% Mis.
GoDaddy	165	42%	Epik	14	3.6%
NameCheap	27	7.1%	FastDomain	10	2.6%
eNom	22	5.7%	Cloudflare	8	2.1%
Network Solutions	20	5.1%	1&1 Ionos	7	1.8%
Tucows	16	4.1%	NameSilo	5	1.3%

Table 5: Top Misinformation Registrars—We show the top ten registars responsible for supporting misinformation websites in our corpus, as well as the proportion of misinformation website each registrar covers.

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 6: Top Advertisers on Misinformation Websites—The distribution of the top 10 advertising trackers found on misinformation websites. Google constitutes three of the 10 advertising domains.

Monetization Strategies

Beyond hosting infrastructure, misinformation sites also rely on online advertising and direct donations to stay online. In this section, we analyze the role monetization providers play in supporting misinformation websites.

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 (Silverman, Singer-Vine, and Thuy Vo 2017; Silverman 2017), 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 DoubleClick is used on 21.6% of all misinformation sites (Table 6).

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 7: 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 also Bonferroni corrected ($n = 132$) and statistically significant.

Compared to all mainstream sites with ads, misinformation sites with ads disproportionately rely on several ad providers, most notably RevContent and DoubleClick (Table 7). 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) and that this difference in proportions is very large. 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 is not prohibited (Silverman 2017). 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" (Love and Cooke 2016). 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 (Graham 2020). All News Pipeline, a conspiracy theory site, laments the decline of ad revenue for itself and other "independent" media sites (Duclos 2018):

> *"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 different ad providers, underscoring the need for coordinated efforts amongst providers for such a mitigation approach. Unfortunately, this does not appear to be happening in practice. Despite RevContent and Google previously claiming to be curbing misinformation on their platforms (Silverman, Singer-Vine, and Thuy Vo 2017; Silverman 2017; Love and Cooke 2016), 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.

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 hate 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, is present in our set of misinformation sites. Re-

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

turn 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"* (Valizadeh 2018).

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"* (DailyStormer 2021).

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.

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%), and Akamai (54%) (Table 8). 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.

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 basis for such decisions. Our results suggest that deplatforming hate speech websites from hosting services has only 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.

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 (Levchenko et al. 2011; K.Sood and J.Enbody 2013; Konte, Perdisci, and Feamster 2015), 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 serving such content.

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 (DDoS-Guard 2021; Cloudflare 2021). It remains unclear whether smaller DDoS protection providers—especially at analogous free tiers—can withstand significant attacks.

Monetization

Most misinformation sites depend on online ads to generate a profit, 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. Although Google is a mainstream ad provider, it is more than twice as common among misinformation than mainstream sites with ads. Extrapolating from anecdotal accounts from smaller hate speech sites, blocking these mainstream monetization channels for sites may be a 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 (Ramachandran and Feamster 2006), we suggest that monetization platforms—both ad providers and payment processors—examine their roles as stakeholders in the online misinformation ecosystem and how they may be in the best position to curb its spread.

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 known problematic 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.

Conclusion

In this paper, we analyze the infrastructure that powers misinformation websites. We show that several providers are disproportionately responsible for serving online misinformation, most prominently Cloudflare, which serves content for 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 rely on monetization platforms like ad networks and donation platforms, and they even rely disproportionately on certain providers like Google and RevContent. Anecdotally, we observe that 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.

Acknowledgments

This work was supported in part by the National Science Foundation under grant #2030859 to the Computing Research Association for the CIFellows Project, NSF Graduate Research Fellowship #DGE-1656518, and the Terman Award at Stanford University.

References

Agarwal, P.; Joglekar, S.; Papadopoulos, P.; Sastry, N.; and Kourtellis, N. 2020. Stop tracking me Bro! Differential Tracking of User Demographics on Hyper-Partisan Websites. In *The World Wide Web Conference*.

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

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

AmmoLand. 2020. Gun Owner Privacy. https://www.ammoland.com/tags/gun-owner-privacy/feed/. Accessed: 2021-01-11.

Arshad, S. 2020. Crawlium. https://github.com/sajjadium/Crawlium. Accessed: 2020-10-08.

Automattic. 2021. Terms of Service. https://wordpress.com/tos/. Accessed: 2021-01-15.

BBC. 2017. Daily Stormer: Cloudflare drops neo-Nazi site. *BBC* URL https://www.bbc.com/news/technology-40960053.

Breland, A. 2017. Alt-right Twitter rival may lose its web domain. https://thehill.com/policy/technology/351169-alt-right-twitter-rival-might-lose-its-domain. Accessed: 2021-01-11.

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

Budak, C. 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, 139–150. New York, NY, USA: Association for Computing Machinery. ISBN 9781450366748. doi: 10.1145/3308558.3313721. URL https://doi.org/10.1145/3308558.3313721.

Burch, S. 2017. Russian Internet Providers Ordered to Say Nyet to Hosting The Daily Stormer. https://www.thewrap.com/russian-internet-boot-daily-stormer/. Accessed: 2021-01-15.

Butkiewicz, M.; Madhyastha, H. V.; and Sekar, V. 2011. Understanding website complexity: measurements, metrics, and implications. In *ACM SIGCOMM: Internet Measurement Conference*.

Cloudflare. 2020. Website and Online Services Terms of Use. https://www.cloudflare.com/website-terms/. Accessed: 2021-01-11.

Cloudflare. 2021. Our Plans. https://www.cloudflare.com/plans/. Accessed: 2021-01-15.

Cox, K. 2021. Parler goes dark, sues Amazon to demand immediate reinstatement. *Ars Technica* URL https://arstechnica.com/tech-policy/2021/01/parler-goes-dark-sues-amazon-to-demand-immediate-reinstatement/.

DailyStormer. 2021. Support the Daily Stormer! https://dailystormer.su/contributions/. Accessed: 2021-01-15.

DDoS-Guard. 2021. Service Plans. https://ddos-guard.net/en/store/web. Accessed: 2021-01-15.

DigitalOcean. 2020. Acceptable Use Policy. https://www.digitalocean.com/legal/acceptable-use-policy/. Accessed: 2021-01-11.

Duclos, S. 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. Accessed: 2021-01-15.

Durumeric, Z.; Wustrow, E.; and Halderman, J. A. 2013. ZMap: Fast Internet-Wide Scanning and its Security Applications. In *USENIX Security Symposium*.

Englehardt, S.; and Narayanan, A. 2016. Online tracking: A 1-million-site measurement and analysis. In *ACM SIGSAC: Conference on Computer and Communications Security)*.

Farsight Security. 2021. Introducing DNSDB 2.0. https://www.farsightsecurity.com/solutions/dnsdb. Accessed: 2021-01-11.

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

GoDaddy. 2020. GoDaddy Legal Agreements and Policies. https://www.godaddy.com/legal/agreements.

Graham, M. 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.

Greenspan, R. E. 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. Accessed: 2021-01-11.

Hanley, H. W.; Kumar, D.; and Durumeric, Z. 2022. No Calm in the Storm: Investigating QAnon Website Relationships. In *AAAI Conference on Web and Social Media*.

Hao, S.; Kantchelian, A.; Miller, B.; Paxson, V.; and Feamster, N. 2016. PREDATOR: Proactive Recognition and

Elimination of Domain Abuse at Time-Of-Registration. In *ACM SIGSAC: Conference on Computer and Communications Security)*.

Hao, S.; Syed, N. A.; Feamster, N.; Gray, A. G.; and Krasser, S. 2009. Detecting Spammers with SNARE: Spatio-temporal Network-level Automatic Reputation Engine. In *USENIX Security Symposium*.

Ho, G.; Cidon, A.; Gavish, L.; Schweighauser, M.; Paxson, V.; Savage, S.; Voelker, G. M.; and Wagner, D. 2019. Detecting and characterizing lateral phishing at scale. In *USENIX Security Symposium*.

Ho, G.; Sharma, A.; Javed, M.; Paxson, V.; and Wagner, D. 2017. Detecting credential spearphishing in enterprise settings. In *USENIX Security Symposium*.

Hounsel, A.; Holland, J.; Kaiser, B.; Borgolte, K.; Feamster, N.; and Mayer, J. 2020. Identifying Disinformation Websites Using Infrastructure Features. In *USENIX Workshop on Free and Open Communications on the Internet*.

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/. Accessed: 2021-01-11.

Internet Engineering Task Force. 2013. Certificate Transparency. https://datatracker.ietf.org/doc/html/rfc6962. Accessed: 2021-01-15.

Jin, Z.; Cao, J.; Jiang, Y.-G.; and Zhang, Y. 2014. News credibility evaluation on microblog with a hierarchical propagation model. In *IEEE International Conference on Data Mining*.

Konte, M.; Perdisci, R.; and Feamster, N. 2015. Aswatch: An as reputation system to expose bulletproof hosting ASes. In *ACM SIGCOMM: Conference on Special Interest Group on Data Communication*.

K.Sood, A.; and J.Enbody, R. 2013. Crimeware-as-a-service—A survey of commoditized crimeware in the underground market. In *International Journal of Critical Infrastructure Protection*.

Kumar, D.; Ma, Z.; Durumeric, Z.; Mirian, A.; Mason, J.; Halderman, J. A.; and Bailey, M. 2017. Security challenges in an increasingly tangled web. In *The World Wide Web Conference*.

Kumar, S.; West, R.; and Leskovec, J. 2016. Disinformation on the web: Impact, characteristics, and detection of wikipedia hoaxes. In *The World Wide Web Conference*.

Levchenko, K.; Pitsillidis, A.; ha Chachra; Enright, B.; Félegyházi, M.; Halvorson, T.; Kanich, C.; Kreibich, C.; Liu, H.; McCoy, D.; Weaver, N.; Paxson, V.; Voelker, G. M.; and Savage, S. 2011. Click Trajectories: End-to-End Analysis of the Spam Value Chain. In *IEEE Symposium on Security and Privacy*.

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

Love, J.; and Cooke, K. 2016. Google, Facebook move to restrict ads on fake news sites. *Reuters* URL https://www.reuters.com/article/us-alphabet-advertising/google-facebook-move-to-restrict-ads-on-fake-news-sites-idUSKBN1392MM.

Nguyen, V.-H.; Sugiyama, K.; Nakov, P.; and Kan, M.-Y. 2020. FANG: Leveraging social context for fake news detection using graph representation. In *ACM SIGIR: Conference on Information & Knowledge Management*.

Nikiforakis, N.; Invernizzi, L.; Kapravelos, A.; Van Acker, S.; Joosen, W.; Kruegel, C.; Piessens, F.; and Vigna, G. 2012. You are what you include: large-scale evaluation of remote javascript inclusions. In *ACM SIGSAC: Conference on Computer and Communications Security)*.

OpenSources. 2017. OpenSources. https://github.com/bigmclargehuge/opensources. Accessed: 2020-10-08.

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

Popat, K.; Mukherjee, S.; Strötgen, J.; and Weikum, G. 2017. Where the truth lies: Explaining the credibility of emerging claims on the web and social media. In *The World Wide Web Conference*.

Ramachandran, A.; and Feamster, N. 2006. Understanding the Network-Level Behavior of Spammers. In *ACM SIGCOMM: Conference on Applications, technologies, architectures, and protocols for computer communications*.

Rashkin, H.; Choi, E.; Jang, J. Y.; Volkova, S.; and Choi, Y. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Conference on Empirical Methods in Natural Language Processing*.

Robertson, A.; and Liptak, A. 2017. Namecheap has taken down Neo-Nazi site Daily Stormer. https://www.theverge.com/2017/8/20/16170370/namecheap-host-take-down-neo-nazi-hate-site-daily-stormer. Accessed: 2021-01-15.

Sharma, K.; Qian, F.; Jiang, H.; Ruchansky, N.; Zhang, M.; and Liu, Y. 2019. Combating Fake News: A Survey on Identification and Mitigation Techniques. *ACM Transactions on Intelligent Systems and Technology* .

Shu, K.; Wang, S.; and Liu, H. 2019. Beyond news contents: The role of social context for fake news detection. In *ACM International Conference on Web Search and Data Mining*.

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

Silverman, C.; Singer-Vine, J.; and Thuy Vo, L. 2017. In Spite Of The Crackdown, Fake News Publishers Are Still Earning Money From Major Ad Networks. *BuzzFeed News* URL https://www.buzzfeednews.com/article/craigsilverman/fake-news-real-ads.

Valizadeh, R. 2018. RETURN OF KINGS IS GOING ON HIATUS. https://www.returnofkings.com/195790/return-of-kings-is-going-on-hiatus. Accessed: 2021-01-15.

WhoTracks.Me. 2021. WhoTracks.Me. https://whotracks.me. Accessed: 2021-01-11.

Wong, J. C. 2019. 8chan: the far-right website linked to the rise in hate crimes. *The Guardian* URL https://www.theguardian.com/technology/2019/aug/04/mass-shootings-el-paso-texas-dayton-ohio-8chan-far-right-website.

Zeng, E.; Kohno, T.; and Roesner, F. 2020. Bad News: Clickbait and Deceptive Ads on News and Misinformation Websites. In *Workshop on Technology and Consumer Protection*.