Before the
FEDERAL TRADE COMMISSION
600 Pennsylvania Ave., NW
Washington, DC 20580

Business Practices of Cloud Computing Providers	No. FTC-2023-0028

Comments of **NICK MERRILL** & **TEJAS N. NARECHANIA**

EXECUTIVE SUMMARY

The phrase "cloud computing" can refer to any one of a wide range of products and services, from software provisioned as an internet-delivered service, to hosted infrastructural solutions. In this comment, we focus specifically on content delivery networks (or CDNs, for short), one sort of an infrastructure-as-a-service.

CDNs are systemically important to the modern internet. Specifically, CDNs offer, among other things, content caching and cybersecurity services. The NBA, for example, may enlist a CDN to more efficiently distribute video highlights. The CDN, in turn, will distribute and locate the NBA's data near users likely to request it (e.g., Boston Celtics-related content in Massachusetts, and Golden State Warriors-related content in California), in order to improve the performance of the NBA's web services. The CDN can also provide protection against cyberattacks by monitoring large-scale traffic patterns across its clients and blocking malicious activity targeting any one (such as, to continue the example, the NBA).

CDNs thus offer significant improvements over prior, more decentralized models of the internet, in which latency was a bigger problem, and certain cyberattacks were more frequent and disruptive. But extreme concentration in the market for CDN services also poses new and unique risks to the internet. Specifically, such concentration may undermine competition among providers of infrastructural services, the security and resiliency of the internet in other ways, and the web's openness as a platform for speech, commerce, and innovation.

Our research, summarized in this response (and attached in full as an Appendix), highlights that CDNs . . .

1. **. . . operate in a highly concentrated market**. Specifically, a single provider, Cloudflare, accounts for nearly 76% of the entire market for CDN services, and three providers account for 89% of that market.

2. **. . . may contribute to outages of global internet services**. Evidence suggests that failures in—and cyberattacks against—CDNs can render vital internet services

unreachable. These risks are exacerbated by the complex interdependencies that characterize the modern internet's software supply chain and the degree of concentration in the market for CDN services. Such concentration essentially gives rise to single points-of-failure on the internet. Even well-protected assets, such as the websites of national governments and major technology companies, have proved vulnerable to failures at CDNs.

3. **. . . use proprietary algorithms, some powered by machine-learning-based "artificial intelligence," to provide their services**. CDNs use artificial intelligence, or AI, to drive caching decisions and to identify and deflect cyberattacks. While these algorithms help CDNs deliver high-quality service, the network effects and opacity associated such algorithms also raises concerns for concentration, transparency, and verifiability.

4. **. . . can—and have exercised the ability to—censor content**. Because CDNs intermediate the relationship between users and the content they request, they are often technically capable of limiting users' access to internet services. Moreover, because CDNs use algorithms to decide which users may access which services, this power is sometimes exercised in ways that seem opaque to users or is poorly disclosed by CDNs.

By raising these issues, we do not discount the important improvements, described above, brought about by CDNs. Rather, we raise these issues for the Commission's attention so that internet users may continue to enjoy the benefits brought about by CDNs while industry and regulatory authorities address these concerns. Specifically, we recommend that the Federal Trade Commission . . .

1. **. . . study the competitive landscape of the market for CDN services.**

2. **. . . evaluate if and how competition may affect systemic risk related to cyberattacks and technical failures.**

3. **. . . consider technical and regulatory solutions that promote an efficient, secure, and open internet.**

4. **. . . collaborate with agencies that have equities over broadband- and cybersecurity-related concerns, including the Federal Communications Commission (FCC) and Cybersecurity and Infrastructure Security Agency (CISA), among others.**

INTRODUCTION

In general, cloud services commoditize computational capacity, such as storage, processing power, or networking capabilities, and offer this capacity to businesses and users over the internet.

Content-delivery networks, or CDNs, offer one specific category of cloud service. Specifically, CDNs offer storage and caching services for internet content, alongside various cybersecurity and internet networking-related solutions. CDNs have grown substantially in recent years; so much so that they are now of extreme, systemic importance to the internet as a whole.

Notably, the market for CDN services is highly concentrated. One dominant provider accounts for over three-quarters of the entire market, and three providers account for about seven-eighths. This concentration is the starting point for a series of concerns about competition, security, and openness on the internet.

MARKET CONCENTRATION OF CDNS

Market Share of Dominant Providers

The market for CDN services is dominated by three major providers: Cloudflare, Fastly, and Amazon account for 89% of the market, and more than 95% of the market for the top 10,000 websites. Of these, Cloudflare is far and away the largest provider, accounting for over 75% of the whole market.

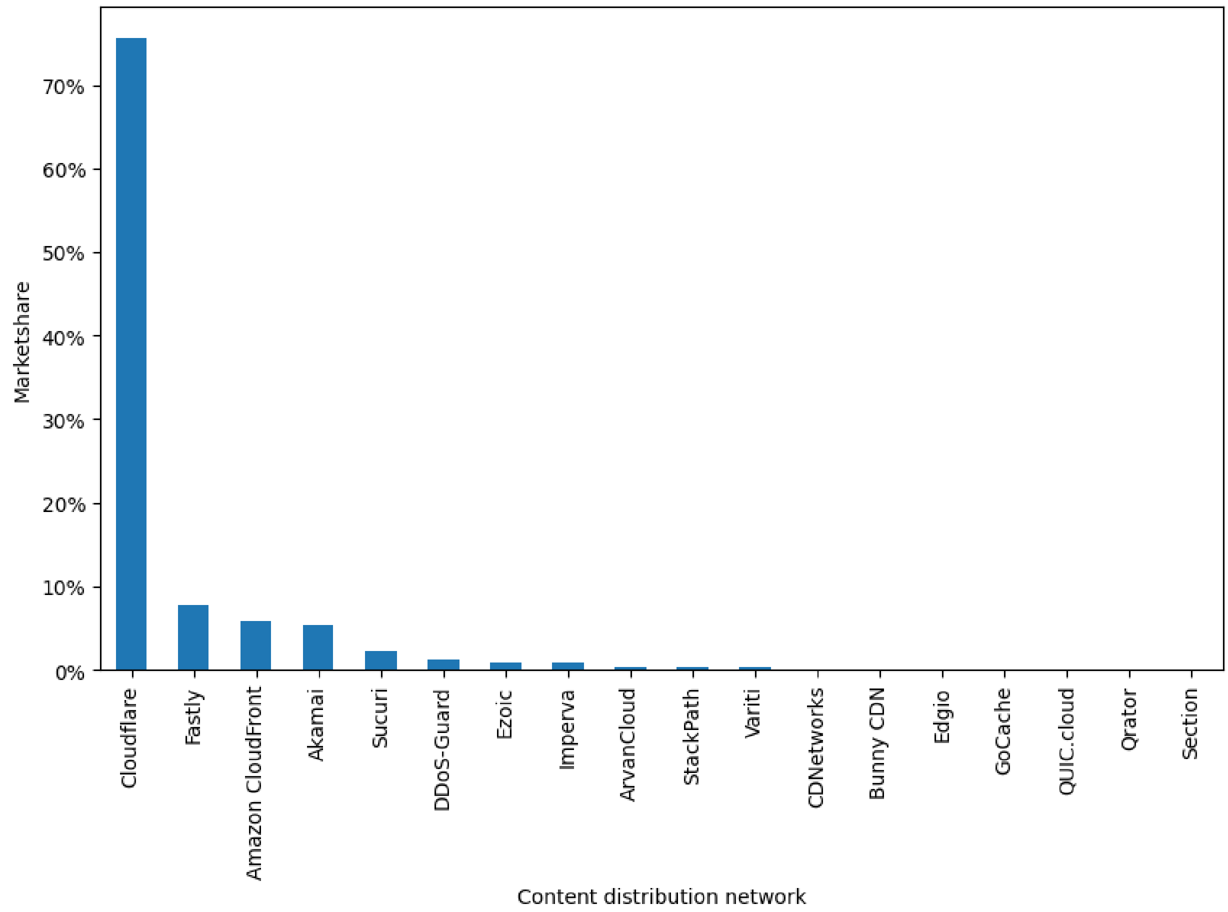

Figure 1. Market Share of CDNs Across All Websites in the Sample Using a CDN to Deliver Service.

We arrive at that figure by beginning with the 15,000,000 most popular websites as measured by Chrome's User Experience Research dataset, the most comprehensive and accurate available dataset of website popularity according to recent internet measurement research.[1] Of those top 15,000,000 websites, an analysis of the packets delivered on behalf of those websites reveals that 23.6% used a CDN to deliver service. We treat these websites as the effective market for CDN services. Of that market, 89% of the packets delivered in response to requests are delivered by one of three CDNs: Cloudflare, Fastly, and Amazon.

This market concentration heightens among the most popular websites. When we focus on the top 10,000 websites in the dataset, we find that 99.9% use a CDN in some capacity. Of the top 10,000 websites, these top three providers, Cloudflare, Fastly, and Amazon, deliver service to

[1] Kimberly Ruth, Aurore Fass, Jonathan Azose, Mark Pearson, Emma Thomas, Caitlin Sadowski, and Zakir Durumeric. 2022. A World Wide View of Browsing the World Wide Web. In ACM Internet Measurement Conference (IMC '22), October 25–27, 2022, Nice, France. ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3517745.3561418

95.6% of the market. 100% of the top 1,000 websites use a CDN, and of those websites, 98.3% use one of these top three providers.

There are some uncertainties associated with our data collection methodology. There is a long tail of less-popular websites—as noted, about three-quarters of all websites—that use no CDN at all. While these account for a large portion of all websites, they account for a vastly smaller portion of web traffic. Also, some providers may use proprietary CDNs that are less easily measured. And other internet traffic, such as for some streaming services or CDN transit, is excluded from or underrepresented in our baseline, given our focus on the most popular websites (rather than the largest consumers of internet bandwidth). But this uncertainty notwithstanding, it remains true that, among websites who make use of a CDN, the market is concentrated: three providers are responsible for 89% of the responses for these websites (and one provider alone is responsible for 75%). Of the top 10,000 websites on the web, almost *all* (over 95%) are reliant on one of three providers.

Our findings echo the results described in other comments to these proceedings. Jonathan Zittrain, for example, filed a co-authored paper regarding the Domain Name System (or DNS). The DNS maps human-readable domain names (e.g., nytimes.com) to machine-readable IP addresses. Although the DNS is itself a decentralized protocol, their research shows that these DNS services are also increasingly centralized—and among the most centralized are Akamai, Amazon Web Services, and Cloudflare. This ostensibly different and ostensibly decentralized component of the internet's core is increasingly consolidated among the same providers. And this is so for reasons similar to the consolidation in the CDN market, as CDN providers include cybersecurity protections (specifically, protections against DDoS attacks) in their DNS service.

The bottom line is clear: The market for CDN services is best characterized as highly concentrated.

Barriers to Entry and Other Causes of Market Concentration

The concentrated nature of the market for CDN services may be the consequence of several interrelated factors, including high barriers to entry and network effects, among others.

Infrastructural Requirements. Because CDNs are infrastructural providers, they require significant investment in infrastructure—global data centers, high-capacity networks, and advanced caching technologies. Moreover, CDNs must build or gain access to data centers located strategically around the world in order to ensure low-latency content delivery. CDNs must also maintain and expand ever-higher-capacity network infrastructures for handling growing volumes of traffic,[2] and navigate partnerships and peering arrangements with internet

[2] Petros Gigis, Matt Calder, Lefteris Manassakis, George Nomikos, Vasileios Kotronis, Xenofontas Dimitropoulos, Ethan Katz-Bassett, and Georgios Smaragdakis. 2021. Seven Years in the Life of Hypergiants' Off-Nets. In ACM SIGCOMM 2021 Conference (SIGCOMM '21), August 23–27, 2021, Virtual Event, USA. ACM, New York, NY, USA, 18 pages. https://doi.org/10.1145/3452296.3472928

service providers (ISPs).[3] All told, developing a CDN is an exceptionally expensive and logistically complex undertaking.

Economies of Scale and Network Effects. These vast infrastructural requirements, together with relatively low marginal costs of service, give established CDNs significant advantages in terms of scale economies.

CDNs also benefit from data network effects. As CDNs grow their customer base, they accumulate valuable information on traffic patterns and user behavior, which can be used to optimize content delivery, enhance security features, and tailor services to certain industries or customer segments. Indeed, CDNs use proprietary algorithms, some powered by machine-learning-based "artificial intelligence" (or AI) to power their services. The data network effects associated with these algorithms give rise to a familiar feedback loop: Entrenched providers leverage their extensive existing data to improve services and solidify their market presence; meanwhile smaller entrants may struggle to gain a foothold in the market.[4] In short, the existing heavyweights in the CDN market have a competitive advantage over new entrants, as their established networks and services provide greater scalability and more comprehensive data.

Lock-In and Switching Costs. Finally, some CDN customers may experience lock-in effects, as migrating to a different CDN provider can be costly and time-consuming, particularly for large enterprises with complex integration requirements. This lock-in effect creates a disincentive for customers to switch providers, making it more difficult for new competitors to enter the market.

CDNs' use of AI, moreover, exacerbates these concerns. As noted, CDNs rely on proprietary algorithms to analyze vast amounts of data generated from user interactions and traffic patterns, using that analysis to make decisions regarding content placement, routing, and security.[5] However, these proprietary algorithms (and the data used to train them) are unique to each CDN provider and their client base, which may create an added layer of dependency for customers. Since these algorithmic optimizations are tailored to a specific CDN's customers, migrating to a different provider may result in reduced performance and security capabilities—at least initially. Cloudflare, for example, boasts "Adaptive DDoS Protection" that is based on having "learn[t] your traffic patterns."[6] Such potential performance degradation is one additional switching cost.

[3] Enric Pujol, Ingmar Poese, Johannes Zerwas, Georgios Smaragdakis, and Anja Feldmann. 2019. Steering Hyper-Giants' Traffic at Scale. In Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies (CoNEXT '19). Association for Computing Machinery, New York, NY, USA, 82–95. https://doi.org/10.1145/3359989.3365430

[4] See, e.g., Tejas N. Narechania, Machine Learning as Natural Monopoly, 107 Iowa L. Rev. 1543, 1584–87 (2022) (describing this feedback loop).

[5] John Graham-Cumming, Bringing AI to the edge with NVIDIA GPUs, The Cloudflare Blog, April 13, 2021, at https://blog.cloudflare.com/workers-ai/

[6] Omer Yoachimik, Introducing Cloudflare Adaptive DDoS Protection—Our New Traffic Profiling System for Mitigating DDoS Attacks, The Cloudflare Blog, Sept. 19, 2022, at https://blog.cloudflare.com/adaptive-ddos-protection/ (boasting "Adaptive DDoS Protection" that "learns your traffic patterns").

EFFECTS OF MARKET CONCENTRATION

The concentration in the market for CDN services has several downstream effects on competition (and competition-related concerns), on cybersecurity, and on internet openness.

Competition

Limited Choice. Most obviously, the degree of concentration in the CDN market limits choice. With a few dominant providers controlling a significant share of the market, customers have only a few options for selecting a CDN that best meets their needs. This lack of choice may cause customers to settle for suboptimal services that do not align with their specific requirements or budget constraints. And the switching costs described above further limit customers' ability to switch providers.

Innovation. Relatedly, the degree of concentration in the market for CDN services may also stagnate the development of new caching or security technologies. As CDNs rest on their inbuilt advantages, resulting from network effects and high switching costs, they may direct less attention to developing new and improved services. In short, a lack of competitive pressure may delay the pace of technological advancement in the industry.

Security

Single Points of Failure. CDNs serve a vital role in the internet's infrastructural ecosystem. And because the vast majority of internet traffic relies on only a few of them, problems at any one can have dramatic effects on the internet as a whole. Consider two recent high-profile incidents. First, in 2019, a software bug at Cloudflare caused a major outage that affected numerous websites and services, including Shopify.[7] The incident demonstrated how a single CDN failure can have widespread consequences across the internet. Likewise, in 2021, a configuration error in Fastly's infrastructure led to a global outage that affected major websites such as Amazon, Reddit, The New York Times, and, perhaps most notably, the website of the United Kingdom's government.[8] The incident underscored the potential risks associated with the centralization of CDN services and the vulnerability of the internet infrastructure to even seemingly minor issues, perhaps in stark contrast to older ideas about the internet's resiliency and its ability to "route around" failures at any one service provider.

The errors and cyberattacks that give rise to such failures can have wide-ranging consequences, as the interconnectedness in the internet's software supply chain risks gives rise to cascading problems. When, for instance, a website's build process relies on certain content hosted by a CDN (an image or a font, say) the failure of that CDN may disrupt that website's functionality—a ripple effect that is felt through the digital supply chain.

[7] John Graham-Cumming, Cloudflare Outage Caused by Bad Software Deploy (Updated), The Cloudflare Blog, July 2, 2019.

[8] Ryan Browne, What is Fastly and Why Did It Just Take a Bunch of Major Websites Offline?, CNBC (June 8, 2021), https://www.cnbc.com/2021/06/08/fastly-outage-internet-what-happened.html.

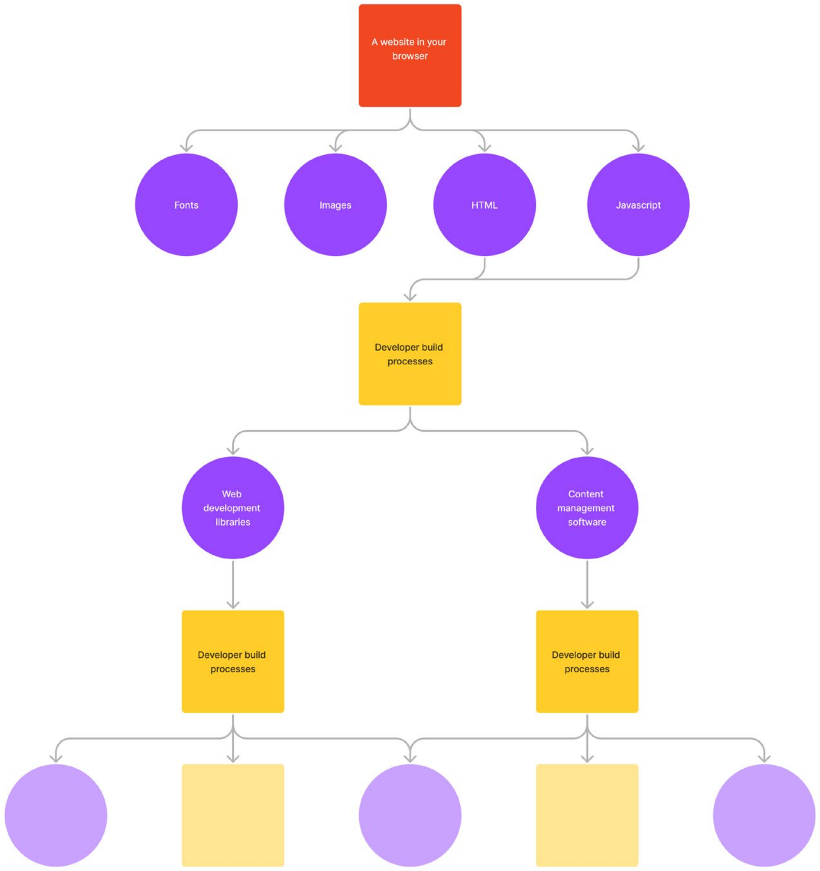
A website in your browser
Fonts
Images
HTML
Javascript
Developer build processes
Web development libraries
Content management software
Developer build processes
Developer build processes

Figure 2. A Schematic Diagram of the Software Supply Chain.

Websites that appear in your browser rely on a variety of static assets (like images and fonts) as well as software assets (like build tools). Those software assets themselves rely on development processes, which themselves rely on software assets, creating a recursive supply chain of "nested" dependencies that can be dozens if not hundreds of layers deep. Assets in purple represent reliances on content that is likely stored with one (or more) CDN provider(s). A failure to deliver any of these assets could cause all downstream products to become unavailable, behave unpredictably, or become impossible to update.

Such problems affect consumers, commerce, and even national security. For one, CDN failures can result in slower page load times, reduced website functionality, and even complete inaccessibility. This degradation in user experience can lead to frustration, decreased customer satisfaction, and potential reputational damage for affected businesses. CDN downtime, moreover, can lead to significant financial losses for businesses that rely on their services. For instance, e-commerce platforms may experience reduced sales and lost revenue due to decreased website availability and slower page load times. And, as CDNs play a critical role in the infrastructure that supports government websites and other vital services, failures or cyberattacks against CDNs can pose risks to national security. Disruptions in critical services may hinder emergency response efforts, compromise sensitive information, or even destabilize essential communication channels during times of crisis.

Verifiability. As noted above, CDNs use proprietary algorithms to deliver cybersecurity services. These algorithms are typically not available for review or analysis by independent experts. Perhaps that is for good reason: Disclosing these measures might make them easier to evade. But, on the other hand, the opacity of the CDNs' underlying methods makes independent, expert verification of their effectiveness (and of the CDNs' claims) challenging.

Openness

Single Points of Control. CDNs intermediate the relationship between users and the content they access online. CDNs cache a wide range of content, such as web pages, images, and videos, in strategically located servers around the globe, on behalf of a wide range of clients. And they deliver this content to a wide range of users worldwide, often without ever needing to use traditional modes of internet "transit" (at least not in direct response to a user request for content). Indeed, CDNs can reach about 76% of internet users by directly sending content to their local internet service provider, without ever needing to enlist the support of a transit provider.[9]

As a consequence, CDNs have direct control over the delivery of content, giving them significant influence over what users see and access on the internet. This influence is exercised in at least two ways.

First, CDNs have the effective power to silence websites, and can do so simply by deciding to cease providing service to a given site. Consider Kiwifarms, an online forum well-known for organizing and supporting vicious, sustained campaigns of stalking and harassment. Nothing we say here is an endorsement of Kiwifarms or the cruelty occasioned by its presence and its members. Cloudflare, in response to public pressure, eventually stopped providing service to the site, a decision which has since subjected Kiwifarms to nearly constant cyberattacks.[10] This reflects the power that a single CDN has to decide which websites are available to users online. While denying service to Kiwifarms seems, in our view, the correct call, we have no guarantees about which entities will, in the future, own a controlling stake in Cloudflare and the causes and websites such entities will and will not support. Indeed, even Cloudflare's current CEO, Matthew Prince, has expressed some discomfort over Cloudflare's vast power over internet

[9] Todd Arnold, Jia He, Weifan Jiang, Matt Calder, Italo Cunha, Vasileios Giotsas, and Ethan Katz-Bassett. 2020. Cloud Provider Connectivity in the Flat Internet. In Proceedings of the ACM Internet Measurement Conference (IMC '20). Association for Computing Machinery, New York, NY, USA, 230–246. https://doi.org/10.1145/3419394.3423613.

[10] Ashley Belanger, Cloudflare Explains Why Kiwi Farms Was Its Most Dangerous Customer Ever, ArsTechnica, Sept. 8, 2022, at https://arstechnica.com/tech-policy/2022/09/cloudflare-explains-why-kiwi-farms-was-its-most-dangerous-customer-ever/

content, suggesting that we would be better served with more democratic control over such decisions.[11]

Second, as noted, CDNs rely on proprietary algorithms to block or screen user requests for content. But these algorithms may rely on data that is biased or unrepresentative, giving rise to content delivery services that discriminate against certain types of content, users, or regions. This results in unfair and unequal access to internet services for some users, and does so in a way that is opaque to users and hardly disclosed by the CDNs themselves.[12]

RECOMMENDATIONS

Competition and Concentration

We recommend that the Federal Trade Commission undertake a closer study of the market for CDN services. Specifically, the FTC should investigate the degree of concentration in the market, including whether any provider has, and has abused, market power. Such a study might encompass an analysis of providers' respective market shares, pricing strategies, and other terms and conditions. Such a study should also assess the effects of concentration on consumer choice, including but not limited to, the availability of alternative providers, the power of customers to negotiate contractual terms with CDNs, and the degree to which CDNs' clients are "locked in" to their existing arrangement. Such a study might also consider the potential benefits and drawbacks of greater competition in the CDN market, including, respectively, increased innovation and greater fragmentation.

Security Risks

We recommend that the FTC also evaluate how concentration in the market for CDN services relates, specifically, to the security and resiliency of the internet's infrastructure. Such a study might encompass an analysis of whether a market characterized by a greater number of providers would offer better security, distribute risk more effectively, or otherwise reduce the potential for large-scale disruptions resulting from attacks on a single provider.

Such a study might also consider the possibility that a more fragmented market might result in less secure and efficient content delivery, or increased complexity in managing and securing the internet's infrastructure. Indeed, given CDNs' reliance on machine-learning algorithms, the CDNs' returns to data may not diminish with scale, given the ever-evolving nature of cyberattacks. Hence, a more fragmented market may result in providers possessing only partial information, undermining the internet's overall security and resiliency.

[11] Matthew Prince, Terminating Service for 8Chan, The Cloudflare Blog, Aug. 4, 2019, at https://blog.cloudflare.com/terminating-service-for-8chan/

[12] Anne Jonas & Jenna Burrell, Friction, Snake Oil, and Weird Countries: Cybersecurity Systems Could Deepen Global Inequality Through Regional Blocking, 6 Big Data & Society 1 (2019), https://doi.org/10.1177/2053951719835238

Technical and Regulatory Solutions

We also recommend that the FTC consider whether any technical or regulatory solutions may help to address any concerns identified in the studies suggested above, and if so, and which ones are likely to be most effective.

For example, the studies we suggest above might conclude that a more competitive landscape would be preferable along some dimensions while also undermining the security benefits that come from the data network effects that inform the CDNs' security algorithms. In such a case, the FTC might encourage the use of federated machine learning or other data-sharing arrangements.

Analogously, the FTC might encourage or support the development of technical standards, akin to the internet's modular and extensible architecture, enabling a more competitive CDN market. Specifically, the Commission might encourage extensible internet architectures that provision caching and security at the network layer, with the effect of structuring a more competitive, and responsive, market.

The FTC might also collate best practices from existing market participants and from participants in adjacent markets, in order to address security risks or to deter risks of private censorship (such as through increased transparency, or the use of independent oversight committees on high-profile "delisting" or "deplatforming" decisions).

Interagency Collaboration

Finally, we acknowledge that other agencies have equities in several of the matters raised above, and so we encourage the Commission to work closely with these other units of the federal government to address complex issues raised by CDNs' modern popularity.

Such agencies include CISA, which is responsible for protecting the nation's critical infrastructure (including infrastructural cloud services providers, such as CDNs) from physical and cyber threats, as well as the FCC, which has long played a significant role in ensuring the competitiveness and openness of the internet and its constituent parts. Other relevant agencies may include the National Institute of Standards and Technology, among several others.

CONCLUSION

CDNs, a systematically important component of modern internet infrastructure, offer significant improvements to data transmission. However, the extreme concentration within the CDN market raises serious concerns regarding competition, security, and online openness. We strongly recommend that the FTC collaborate with other relevant agencies to examine this vital industry, establishing and implementing necessary rules, standards, and best practices to ensure a resilient and competitive CDN landscape.

INTEREST OF COMMENTERS

Nick Merrill is a Research Fellow at the University of California, Berkeley Center for Long-Term Cybersecurity. Tejas N. Narechania is the Robert and Nanci Corson Assistant Professor of Law at the University of California, Berkeley, School of Law. (Our institutional affiliations are noted for identification purposes only. Our comments reflect only our own views, and should not be understood to represent the views of our employer or other colleagues.)

We research and write about matters of internet and telecommunications technology, law, and policy. Our attention to this proceeding is informed by our academic work, and our interest in the sound development of regulatory measures regarding the internet. We have no financial interest in the outcome of this proceeding.

APPENDIX

Nick Merrill & Tejas N. Narechania, *Inside the Internet*, DUKE L.J. ONLINE (forthcoming 2023)