Independent Submission
Internet Draft
Intended status: Informational

D. Lazanski
Last Press Label
M. McFadden
Internet policy advisors, ltd

Expires: April 4, 2024

October 4, 2023

On the Effects of Internet Consolidation
draft-mcfadden-cnsldtn-effects-01

Status of this Memo

Copyright Notice

Abstract

This document contributes to the continuing discussion on Internet consolidation. Over the last several years there have been many types of discussions around consolidation at a technical level, an economic or market level and also at an engineering level. This document aims to discuss recent areas of Internet consolidation and provide some suggestions for advancing the discussion.

Table of Contents

1. Introduction

The origins of the Internet was and continues to be decentralised. Resilience, security and best-effort delivery of data and information

on all layers of the Internet works best in a decentralised manner. But over the last several years there have been discussions on how the Internet is becoming "centralised"¥ or "consolidated"¥ (see section 2, below).

Internet consolidation is "the process of increasing control over internet infrastructure and services by a small set of organizations." [1] Let us consider two general categories of concentration: "player" and "layer". By player concentration, we mean the aggregating of a market to a small number of providers for a particular service. Layer concentration means the combining of functions within a given layer. An example of "player" concentration would be a relatively small number of email service providers who offer billions of users email service. [2] Or the number of web search providers or even web browser offerings. [3]

As defined in draft-nottingham-avoiding-Internet-centralization [4] "centralization" as the ability of a single entity or a small group of them to exclusively observe, capture, control, or extract rent from the operation or use of an Internet function. Furthermore, "centralisation"¥ as noted in the Internet of three Protocols is that one or two or three single protocols are being used for everything rather than one protocol for one operation as is a guiding principle of protocol design until now.

The Internet is being centralised and, thus, consolidated on all layers of the Internet and it is essential to recognise the technical, political and economic reasons for this happening. The rest of this draft will focus on different aspects of the issue of consolidation.

2. Acknowledgement of Other Drafts on This Topic

This document recognizes that the topics of protocol design and centralization have been addressed by several people. In this section we take a moment to recognize that we are not the first to come to this topic, nor will we be the last. Our purpose is to examine the forms of centralization and how protocol design impacts them.

In section 1 above we cited the draft from Mark Nottingham that discusses centralization in Internet protocols and relates it to consolidation of power.[5] The draft goes on to identify possible reactions to centralization and specifically what Internet protocols should do to limit or mitigate centralization.

Another draft (now expired) explored a slightly different angle. In draft-arkko-iab-internet-consolidation-02, the authors consider the topic from the perspective of how available technology and Internet architecture drives different market directions.[6] This draft ends with a call to action that emphasizes open interfaces, specific standardization choices and the benefits of open source development and the need for further research.

Another important contribution to the discussion is a paper "Centrality and the Internet"¥ published by Geoff Huston on his blog.[7] The paper explores the historical precedents for consolidation and the consequences of having large organizations control important parts of specific sectors of the economy. It finishes with a look at the role of regulation in ensuring that the market functions properly and the impact of advertiser funding in creating a small number of dominating incumbents.

Another recent paper, by a team from The Netherlands and Brazil, examines consolidation in the hosting industry.[8] The paper focuses on that industry and shows how it is heavily concentrated: 10 hosting providers account for most of the hosting for all TLDs considered. While European ccTLDs have a strong hosting industry, US-based providers have been continuously conquering the market, especially in the high end of it - the popular domain names, which poses challenges for the European Union's goals of digital sovereignty.

Both the Internet Society and participants of the IETF have published on the subject of consolidation in 2019. At the IAB's Design Expectations vs. Deployment Reality in Protocol Development Workshop **2019 a handful of the participants discussed concentration and** consolidation. Jari Arkko discussed the impacts of consolidation on the Internet infrastructure in a document for the IETF[9], with the document identifying issues including loss of resilience and increased risk of surveillance. It goes on to note that "it seems prudent to recommend that whenever it comes to Internet infrastructure services, centralised designs should be avoided where possible".[10] From networks to applications, the overarching theme was that consolidation is taking place from one end of the Internet to the other.

Additionally, the Journal of Cyber Policy published a special edition on Consolidation of the Internet. Topics in this special issue included market concentration and security, DNS consolidation, supply chains, interoperability and Internet architecture. However, much is still yet to be discussed on consolidation at most layers of the Internet stack. [11]

The discussion of consolidation primarily focuses on Internet services and data. However, it is important to draw attention to the issues and risks of consolidation at other layers of the Internet beyond just the application layer. The application layer is directly user-facing and, as a result, is what users experience. But the underlying infrastructure and protocols are also going through consolidation as they develop. The complete end-to-end encryption model forces data into endpoints which consolidates data into and handful of companies. Furthermore, protocol standards are facilitating this consolidation.

3. Background to Consolidation Issues and the Role of Standards

The Internet is being consolidated at all layers, from the application layer to the network layer. In the context of search online Google has 84% of all searches online.[12] But market consolidation is not limited to the Internet. It happens when economies of scale provide highly aggregated firms an advantage. For the last three decades, we have witnessed concentration occurring not only in telecommunications, but in the financial sector as well. Concern is growing over the fact that financial institutions are only using cloud services from a handful of cloud service providers.[13] The acceleration of consolidation has been assisted by cloud technologies, such as occurred with email. Thanks to ease of use enabled by cloud hosting, services like email and online payments can be accessed via a web browser.

In other market consolidation cases, fewer Internet standards are in play. In the case of home assistant tools such as the Amazon Echo or Google Home Assistant, communication from these devices to their respective clouds is largely proprietary in nature. In particular, the information models and schemas they use are not exposed to the outside world. This is because the bulk of the service is performed by the cloud, with relatively little processing occurring in the home. This two-sided model eliminates the lengthy standards development process, thereby permitting faster service improvements.

On the Internet over previous decades, numerous Internet Service Provider (ISP) markets were subject to deregulation, disaggregation of customers by regulatory requirement, consolidation, and to some extent, re-regulation.

In years past, standards have been viewed as a means to prevent barriers to entry. During the 1980s, AT&T was required to abide by standards as part of the consent decree that resolved antitrust litigation, leading to the ability of anyone to connect a telephone to its network. By 1994 standards were recognized as a means to prevent technical barriers to trade (TBT) during the Uruguay Round of the World Trade Organization.

The QUIC protocol[14] is an example of the consolidation between layers of the Internet - and not at the application layer. Designed and deployed as a transport layer protocol, it effectively replaces TCP at the network layer while also adding improved security. The result is the merging or consolidation of three layers. QUIC should improve efficiency and delivery of applications, but also forces all data to be managed at the endpoint, which in this case is a browser, making it more difficult to manage traffic at the network layer.

4. Overarching Issues Related to Consolidation

4.1. Essential Taxonomy of Internet Consolidation

Discussions at the IETF (and elsewhere) have shown that different people have different views of how consolidation expresses itself. While there is little argument that the increasing control of Internet infrastructure and services is being coalesced into the hands of a small number of organizations. However, that consolidation expresses itself in a variety of ways.

Another draft suggests a potential taxonomy of consolidation and proposes four main categories: [15]

- Economic consolidation

- Traffic and infrastructure consolidation

- Architectural consolidation

- Service and Application Consolidation

4.2. Technical

Consolidation has led to the development of a few, large Internet companies which consumers are using by way of platform consolidation, as mentioned above. But consolidation also has led to the development of protocols which are developed and used by these few, large Internet companies to control traffic flow and data capture as well.

Overarching technical issues related to consolidation include an over-reliance on one or two entities and a handful of protocols. Certain stakeholders who have developed and implemented these protocols manage the updated and upgraded versions of the protocols.

"Did the IETF create a better internet when it approved DoH?" There's a lot of disagreement about that, but what has upset many is that DoH was a surprise - the IETF standardised it without consulting some who it

was likely to affect," it says in RFC 8890 [16] However, there was little multistakeholder consultation and discussion prior to the adoption of DoH. This was more of a rapid development and deployment process, without the market driving the use cases and uptake. By forcing the concentration of the data at the endpoint, the data is consolidated into the service provider at that endpoint.

4.3. Economic

According to the Internet Society's 2019 report Consolidation In the Internet Economy the Internet economy is broadly defined as, "economic activities that either support the Internet or are fundamentally dependent on the Internet's existence."[17] Internet applications, service infrastructure and access provision are the primary three areas of economic activities on the Internet.

One focus of consolidation is around the concentration of power - consumer, technical and financial - into a handful of large Internet companies. The first point of engagement with any of these companies, including Facebook and Google, is through consumer applications. The ability to easily understand consolidation at an application layer, because of the widespread and common use of Facebook and Google, has caused the focus of consolidation and anti-competitive issues from policymakers and politicians to be at the application layer.

However, consolidation doesn't always have its downsides. Consolidation allows for economies of scale, investment in infrastructure and the ability for small and medium enterprises to buy and use services, like cloud storage, content distribution networks and security technology, without having to build them from the ground up every time. However, the lack of market diversity means a lack of competition which, in turn means a lack of innovation and a lack of consumer choice.

Amazon offers affordable cloud services and Cloudflare is one of only a handful of companies that are content delivery networks at a large scale. So large, in fact, that a substantial amount of Internet traffic transits through Cloudflare's servers, though there are many thousands of small CDNs. Rather than each and every Internet application company create their own storage and content delivery network, it is easier and more affordable to outsource both to other companies. Because of the cost of CDNs at scale, few companies offer these services.

The market should be a regulating factor in consolidation. New entrants and competition in a market creates options for consumers that potentially pulls them away from popular websites and applications. When a market is not competitive or viable, regulation and anti-trust measures can intervene to remedy a consolidated market which is tending

towards or has achieved monopoly status. Legal and regulatory intervention, however, tends to create its own set of issues as seen through several decades of EU intervention in big tech starting with Microsoft in 2004. Unintended consequences with regulatory or legal intervention may skew the market even further.

Economics is driving protocol design in a couple of different ways. First, participation in standardization is open and free, at least in one sense and for the IETF. However, attending the IETF in person requires a financial commitment - not just for registration, but the travel, hotel and expenses are costly. The very organizations that can afford to attend in person are the ones facilitating consolidation.

4.4. Security

Consolidation of protocol development which has facilitated the secure, end-to-end encryption of information going over networks in recent years. New technologies such as DNS-over-HTTPS (DoH) and DNS-over-TLS (DoT) standardised through the IETF process allow for confidential look-up of DNS queries. However, it has required updates to many DNS servers and operating systems. The implementation of this protocol enables circumvention of DNS filtering which ISPs offer for protection from malicious websites and software on the network.

This is a form of market consolidation based on development choices by several large companies. These development choices are often technically opaque without transparency of what happens when updates take place, resulting in more difficulty when trying to troubleshoot security issues.

The development of these protocols, while providing increased privacy and addressing issues concerning government surveillance, have for another unintended consequences which is promoting consolidation.

Consequences of the security of the global Internet are evident. On June 8, 2021, a global outage of Fastly, a content delivery network (CDN), was caused by a software update which included an undiscovered bug. [18] While this was resolved within a working day, one of the main causes of the outage was a dependency on the limited number of CDNs running services in the cloud. Other CDNs, which resolved traffic via Fastly for redundancy, were also taken down as a result of the Fastly outage. This dependency is caused by consolidation and a concentration of infrastructure. A highly consolidated CDN network facilitates a less secure environment because of the weakening of resilience. [19]

5. Centralization versus Consolidation

The words 'centralization' and 'consolidation' are often used interchangeably when discussing the idea of concentration within the internet. However, centralization and consolidation are, in fact, different. Consolidation is an economic choice one that is driven by economies of scale and efficiencies of work. Consolidation through economic choices causes the outcome to be a centralized way of building Internet architecture and, thus, a centralized market with limited choices of technical and service options.

Another draft [20] carefully considers the distinction between centralization and consolidation and concludes that decentralized technology - by itself - does not guarantee decentralized outcomes. That same draft describes consolidation as "the ability of a single entity or a small group of them to exclusively observe, capture, control, or extract rent from the operation or use of an Internet function."¥ That draft is careful to identify "Centralization"¥ as the source of consolidation.

6. Implications of Consolidation on Internet Architecture

6.1. The Changing Architecture of the Internet

The phenomenon of consolidation may be in the eyes of the beholder. A government may see market failure or a need for regulation. [18] A civil society advocate may see it from the point of view of privacy or free speech . For the purposes of this draft we view it from the perspective of the underlying architecture of the public Internet.

Consolidation in the Internet's architecture is not a new development. The approach of providing intermediaries to deliver service or content rather than the more traditional end-to-end approach has been in place for more than a decade. However, it is possible to argue that the architecture of the Internet has changed dramatically in the last decade.

The architecture of the Internet is always changing. New services, applications and content mean that the market creates new ways to deliver them. Consolidation clearly has economic, social and policy issues, but it is important to understand how consolidation affects the underlying architecture of the Internet. The impact of intermediaries on architecture is often not obvious.

The use of intermediaries in the Internet's architecture may include the use of third parties to provide services, applications or content. In the early days of the Web, this was evident when rendering a web

page that included content from multiple sources. In today's Internet the intermediaries are not so obvious. Authentication servers, content distribution networks, certificate authorities, malicious content protection and DNS resolution services are all examples of tools provided to the Internet by intermediaries - often without the knowledge or approval of both endpoints.

Having intermediaries embedded in the architecture is a different effect from having them embedded in the service structure. The domination by a few companies of the content and application layer is largely an economic effect of scale. On the other hand, there is a prevalent belief that the Internet puts intelligence at the edge. While that may have been true in the past, it is hard to argue that this is a feature of the contemporary Internet.

There is a suggestion that the network simply provides for the transport of data. There are almost no network connections like that in today's Internet. A consumer's view of the Internet is limited by unseen intermediaries of many types - some delivering positive services, others not. In either case, a consumer on the Internet seldom makes choices about those intermediaries: they are simply part of the fabric that makes up the Internet.

It is into just consolidation from the perspective of a consumer. Almost all important parts of the architecture have been affected by consolidation: DNS resolution, access service, transit provision, content distribution and authorization. Consolidation in these areas has a direct effect on engineering and protocol design.

6.2. The End-to-End Principle Redux

The end-to-end principle is the idea that reliability and trustworthiness reside at the end nodes of networks rather than in the network itself. In other words, the idea was that the network itself was dumb and intelligence was at the edge or end. However, Internet architecture is evolving in such a way that this principle is changing.

Networks and the devices on the networks are acting as access consolidators. While, in the past, the network was a simple transporter of bits, today's networks see intermediaries consolidating both access and the delivery of information (e.g. streaming media). For example, 5G will allow for different services, systems and use cases at a very specific level. Network slicing in 5G will concentrate services like video on demand into concentrated - and consolidation - areas on a network. [21] In other words, as specific types of services are relegated to a segregated part of a network, the availability and access of that service is limited to accessing a specific network.

Depending on the type of device or maturity of the network infrastructure available at the point of the attempted access, options for access might be limited. If a network slice on 5G is where a specific service is located, for example, but it is only possible to use a 3G mobile network, then the service is unavailable. Thus, the service is only available on a consolidated part of the mobile network. Another change is how the layers of the Internet, as discussed in the QUIC example, are consolidating. Differentiation among layers is fading fast with the development of applications which require network access and control.

Rapidly, the end-to-end principle is becoming the edge-to-edge principle. The layers of the internet are morphing into several consolidated layers and it is becoming difficult to differentiate between the end or edge, and also nearly impossible to ensure the reliability of the internet because of it. But the important part of this is the network is not dumb. Data processing, storage and highly evolved services (including custom data and metadata processing at the edge) means that the 'dumb' network is no longer dumb.

If the number of organizations that provide those "network services" that we rely upon is small, our dependence is higher. In extreme cases of engineering, we put ourselves at risk of engineering a single point of failure. But also if organisations can't and won't enter the market, the market is left with very few options and choices.

The trend toward highly specific and concentrated processing, as well as the drive for highly customised applications and services will drive the Internet away from an end-to-end principle. This will create not a network of networks, but a mesh. If the mesh is dependent on a small number of very large providers through consolidation, we will have engineered a single source of failure into the Internet.

7. Intermediaries and Consolidation

Internet privacy concerns have encouraged protocol designers to take a more aggressive approach to ensuring privacy in communications. In the past, a secure channel using technologies such as TLS or IPsec provided a way to ensure that point-to-point communications was protected while information was in transit. Providing privacy (and authentication of the data stream) occurred between the endpoints of communication.

However, it became widely recognized that this was insufficient. In particular, a secure channel between two endpoints does not guarantee that the information will remain private at the endpoints. As the importance of privacy increased, so too did the attempt to fashion protocols that increased the protection of the data a the endpoints.

A draft from the IAB describes the technique for separating the data and metadata visible to diverse parties in network communication as "privacy partitioning."¥[22] It notes that a group of IETF working groups are using this intermediary strategy as a protocol-based, technical approach to improving privacy at the ends of network connections. The working groups involved include OHAI, MASQUE, Privacy Pass and PPM. All four have in common a general strategy of using an intermediary to provide a higher level of privacy for endpoints.

The use of intermediaries is nothing new: we have had HTTP proxy services in the Internet almost since the advent of the Web. What has changed is the dominance of privacy preservation in protocol design. The intermediaries that provide the privacy partitions are in a special and notable place in a network connection. The former end-to-end principle drops away and in its place are two connections: one between an end user and the intermediary and the other between the intermediary and the requested service or application.

The risk of consolidation to the is approach would mean that a dominant set of large companies provide the intermediary services. That would lead to the possibility of collusion with the consequence that no privacy was actually provided. A centralized service providing the "privacy partitioning"¥ could log requests and share information about patterns of use or actual, specific user information.

The result is that "privacy partitioning"¥ needs to be considered as part of the consolidation landscape. The result of having a very small number of dominant providers acting as the intermediaries would lead to some of the same risks as economic or traffic consolidation already exhibit.

8. Implications of Consolidation on Protocol Design

8.1. Does Protocol Design Really Affect Consolidation?

As noted in "Internet of Three Protocols"¥draft, "One of the guiding principles of designing a protocol in the original Internet community was "the protocol is not complete when everything possible has been added, but rather when everything possible has been removed." This is so that security, scalability, resilience and observability can be ensured. However, the recent trend has been towards having a few protocols, but having those protocols do all things.

Though Internet protocol development should be multistakeholder, but standards development is subject to vested interests, personal approaches and commercial realities.[23] Developing protocols, and standards more generally, takes time, much discussion and a bottom-up

approach. However, commercial organizations have different goals in the process of trying to standardize protocols. Larger organizations have more resources dedicated to protocol and standards development. Larger organizations with staff specifically dedicated to standards tend to have the ability to push for their proposals and their protocols. There is no coincidence that these companies are the ones that have facilitated consolidation on a commercial level and are facilitating consolidation on a protocol level.

8.2. Case Studies in Consolidation and Protocol Design

8.2.1. DNS over HTTPS (DoH)

The development of encrypted DNS, specifically DNS-over-HTTPS (DoH), has been driven by a desire to show full end-to-end encryption of network connections. The protocol was completed and the DoH working group wound up in March 2020 despite the absence of both resolver discovery and selection mechanisms. This may be addressed in the future.[24]

Client software is developing with interim discovery solutions which almost always favour the large, cloud-based resolver operators. This is leading to a situation where users are being presented with a very small number of pre-configured resolver options irrespective of their location - in some client software as few as three or four options may be presented. [25] Currently, there are many thousands of servers operating without DoH.

It is likely that most of the DNS traffic will be consolidated onto a handful of global operators, if multiple options for discovery mechanisms are not developed. The impact that such a loss of diversity of providers may have on the long-term resilience of DNS should not be underestimated. [26] Nor should the attractiveness of these potential network chokepoints to attack be overlooked either to access consolidated data or launch an attack from. One danger is that if DNS traffic is concentrated onto a small handful of global operators and turned 'automatically-on' the result would be default adoption by the vast majority of the Internet's clients. The suggestion that there were mechanisms for users to opt-out would not matter in the face of statistics that regularly show that users almost never change default settings. Currently, the deployment approach for DoH is opt-in. For CDNs, DoH default-on would disrupt and render CDN geolocation designed to manage traffic flows more efficient closer to the desired delivery location. Thus, protocol design decisions that are enshrined in default settings will become the norm. In this case, default on, which facilitates consolidation, will become standard.

By routing the DNS over HTTPS, it becomes much easier to track user activity through the use of cookies. Therefore, a protocol that was developed to enhance user privacy and security could actually undermine both: privacy through the use of cookies and security by consolidating DNS traffic onto far fewer resolver operators that are far more attractive targets for malicious actors of various types.

8.2.2. Encrypted Server Name Indication (eSNI)

Options to encrypt the Server Name Indication (SNI) have been explored in the TLS working group but to date it has not been possible to develop a solution without shortcomings. This flaw in the encrypted SNI (eSNI) options under evaluation required a rethink in the approach being taken.

The solution now proposed, Encrypted Client Hello (ECH, previously called ECHO) assumes that private origins will co-locate with or hide behind a provider (CDN, application server etc.) which can protect SNIs for all of the domains that it hosts.[27] Whilst there is logic in this approach, the consequence is that the would-be standard encourages further consolidation of data to aid privacy. What it does not appear to consider is the attractiveness of this larger data pool to an attacker, compared with more dispersed solutions.

8.2.3. Oblivious HTTP

Oblivious HTTP (OHTTP)[28] is a relay based intermediary system that attempts to provide an extra layer of privacy by incorporating per-message encryption in the relay exchange. A client sends a request to an Oblivious Relay which is not allowed to read its contents. The request is forwarded to an Oblivious Gateway which is able to decrypt the messages but does not know the identity of the client or any metadata (for instance, source IP address) related to the client.

The key to OHTTP's privacy features is that the client metadata and request data are separated into separate contexts: the goal is that no entity (other than the client) can see both contexts.

The major risk in OHTTP is collusion across those contexts. If a small number of providers of the OHTTP services dominated, the risks of collusion might be expanded - specifically, protections against collusion and the exposure of user identifying information would be greater in a marketplace without a variety of servers to provide the service.

9. Potential Technical Risks

There are a number of potential risks to the security, stability and performance of the Internet and many of them are well articulated in draft-livingood-doh-implementation-risks-issues-04 [29], but some notable ones are:

1. **Significant operational shift of the global Internet from a highly** distributed to a centralised system. This would impact both security and resilience.

2. **Decreased stability due to the fact that a centralised system will** have higher fragility, fewer points of failure and greater impact on the system when it does fail.

3. **Increased security issues caused by the reduction of number of** recursive DNS operators. [see https://hbswk.hbs.edu/item/evidence-of-decreasing-internet-entropy-the-lack-of-redundancy-in-dns-resolution-by-major-websites-and-services][30] Lack of distributed and recursive DNS creates a lack of redundancy for when security attacks hit parts of the Internet.

4. **Loss of security threat visibility due to degraded ability to use** DNS blocklists and overall network management for malware, phishing, spam, DDoS and etc if DNS management is consolidated into a few operators.

5. **Reduced diversity in the Internet ecosystem. Diversity creates** greater redundancy, resilience and agility to respond to attacks, outages and network issues.

10. Metrics for Consolidation

It is completely natural, when thinking about how to address the problem of consolidation, to ask how to measure it. If a set of metrics were agreed for consolidation, then those metrics could be assessed at a specific point in time as well as examined in time-series. Recent research shows that this may be possible.

In recent work [31], the Internet Society is able to visualize data on the distribution of market shares of core Internet services and infrastructures. The goal was to find metrics that would support the evaluation of how services were concentrated among a small set of actors - or, a small set of countries. By establishing the metrics, the researchers could also track how this concentration changes over time.

The research looks at two different views of market concentration:

- Market concentration - defined as the concentration of providers in a given market; and,
- Country market shares - defined as the jurisdiction of providers in a given market.

The researchers then calculated two separate values to determine concentration of service provision.

The first is the Gini Coefficient. According to the researchers, "The Gini coefficient measures the degree of inequality in a distribution and is widely used in economics to measure wealth and income inequality. It ranks income distribution on a scale between 0 and 1, where 1 means complete inequality (one actor owns all the shares) and 0 means perfect equality (every actor has the same share)."¥

The second is the Herfindahl-Hirschman Index. Again, according to the researchers, "The Herfindahl-Hirschman Index (HHI) is a commonly accepted measure of market concentration and is calculated by squaring the market share of each firm competing in a market, and then summing the resulting numbers. HHI values are in the range 0 to 10,000. HHI values over 2,500 indicate highly concentrated markets."¥

Content Delivery Networks (CDNs) ranked very high in a recent sampling of the data. The Gini Coefficient of 0.86 is dramatically close to 1.0 and represents a finding of significant market concentration. All the other services surveyed by the researchers had metrics greater than 0.50:

- Top Level Domains - 0.77
- DNS Servers - 0.72
- Data Centers - 0.67
- SSL Certificates - 0.66
- Web Hosting - 0.64

Recalling that Herfindahl-Hirschman Index values over 2,500 indicate highly concentrated markets, CDNs once again ranked highest for market concentration (5,924). Only SSL Certificates ranked above 2,500 among the other indicators.

11. Security Considerations

While this document does not describe a specific protocol, it does discuss the evolving architecture of the Internet. Changes to the Internet's architecture have direct and indirect implications for the Internet's threat model.

Specifically, the changes to the end-to-end model (see section 4.2 above) have inserted new interfaces which must be reflected in security considerations for new protocols.

12. IANA Considerations

This document requests no actions on the part of IANA.

13. Conclusions

This document seeks to rekindle and restart the discussion on consolidation. As argued above, Internet consolidation is happening at different places and different layers of the Internet. Though there has been interest in the Internet consolidation in the past, now is the time to start the discussions again.

14. References

14.1. Informative References

[1] Considerations on Internet Consolidation and the Internet Architecture [draft-arkko-iab-internet-consolidation-02]. Expired

[2] As of April 2022, Apple has over 57% of the email client market, including on mobile devices, and Gmail accounts for over 29%. 85% of the global email client market is made up of only two clients. https://www.litmus.com/blog/email-client-market-share-april-2022/

[3] Google has over 84% worldwide search market share as of January 2023 https://www.statista.com/statistics/216573/worldwide-market-share-of-search-engines/

[4] Internet Centralization: What Can Standards Do? [draft-nottingham-avoiding-internet-centralization-09]. March 2023.

[5] Ibid. See [4] above.

[6] Ibid. See [1] above.

[7] Centrality and the Internet, https://www.potaroo.net/ispcol/2021-06/centrality.html April 2021.

[8] Hosting Industry Centralization and Consolidation, L. Zembruzki, R. Sommese, L. Z. Granville, A. Selle Jacobs, M. Jonker and G. C. M. Moura, NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium, Budapest, Hungary, 2022. https://ieeexplore.ieee.org/abstract/document/9789881/

[9] Design Expectations vs. Deployment Reality in Protocol Development Workshop 2019, Internet Architecture Board https://www.iab.org/activities/workshops/dedr-workshop/position-papers/

[10] Centralised Architecture in Internet Infrastructure [draft-arkko-arch-infrastructure-centralisation-00].

[11] IBID page 5.

[12] Browser & Platform Market Share January 2021 https://www.w3counter.com/globalstats.php

[13] Cloud providers pose potential risk to banking sector: Treasury report https://www.bankingdive.com/news/cloud-providers-pose-potential-risk-banking-sector-treasury-report/642428/

[14] RFC 9000, QUIC: A UDP-based Multiplexed and Secure Transport, https://datatracker.ietf.org/doc/rfc9000/

[15] A Taxonomy of Internet Consolidation, https://datatracker.ietf.org/doc/draft-mcfadden-consolidation-taxonomy/

[16] RFC 8890, The Internet is for End Users. Nottingham, Mark. August 2020. https://www.rfc-editor.org/info/rfc8890

[17] Consolidation In the Internet Economy, Internet Society, 2019.https://future.internetsociety.org/2019/consolidation-in-the-internet-economy

[18] Fastly Blog, June 8, 2021. https://www.fastly.com/blog/summary-of-june-8-outage

[19] The Deeper Root Cause of the Fastly and Akamai Outages, CircleID, June 28, 2021. https://www.circleid.com/posts/20210628-the-deeper-root-cause-of-the-fastly-and-akamai-outages/

[20] Ibid., see [4] above.

[21] See Google, antitrust and how to best regulate big tech, The Economist, 7 October 2020 https://www.economist.com/business/2020/10/07/google-antitrust-and-how-best-to-regulate-big-tech

[22] Partitioning as an Architecture for Privacy, https://datatracker.ietf.org/doc/draft-iab-privacy-partitioning/

[23] Dominique Lazanski, Governance in international technical standards-making: a tripartite model, Journal of Cyber Policy, 4:3, 362-379, 2019.
https://www.tandfonline.com/doi/full/10.1080/23738871.2019.169.6851

[24] DNS over HTTPS (doh)
https://datatracker.ietf.org/group/doh/about/

[25] At the time of writing, the Firefox browser presents a list of three pre-configured resolver options to North American users: Cloudflare, NextDNS and Comcast.

[26] Cloudflare DNS goes down taking a large piece of the Internet with it, 17 July 2020. https://techcrunch.com/2020/07/17/cloudflare-dns-goes-down-taking-a-large-piece-of-the-internet-with-it/

[27] TLS Encrypted Client Hello draft-ietf-tls-esni-07
https://tools.ietf.org/html/draft-ietf-tls-esni-07

[28] Oblivious HTTP, https://datatracker.ietf.org/doc/draft-ietf-ohai-ohttp/ but also see the same approached used for the DNS: RFC 9230, Oblivious DNS over HTTPS, https://datatracker.ietf.org/doc/rfc9230/

[29] Centralized DNS over HTTPS (DoH) Implementation Issues and Risks, https://datatracker.ietf.org/doc/draft-livingood-doh-implementation-risks-issues/ (expired draft).

[30] Evidence of Decreasing Internet Entropy: The Lack of Redundancy in DNS Resolution by Major Websites and Services, Samantha Bates, John Bowers, Shane Greenstein, Jordi Weinstock, and Jonathan Zittrain, https://hbswk.hbs.edu/item/evidence-of-decreasing-internet-entropy-the-lack-of-redundancy-in-dns-resolution-by-major-websites-and-services March 2018

[31] Internet Society Pulse - Market Concentration, Internet Society, https://pulse.internetsociety.org/concentration, September 2023

15. Acknowledgments

Many thanks to all who discussed this with us, especially Jason Livingood, Geoff Huston and Jari Arkko.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Dominique Lazanski

Last Press Label

London, UK

Email: dml@lastpresslabel.com

Mark McFadden

Internet policy advisors ltd

Chepstow, Wales, UK

Email: mark@internetpolicyadvisors.com