

Workgroup: Internet Engineering Task Force

Internet-Draft:

draft-mcfadden-consolidation-taxonomy-03

Published: 23 October 2023

Intended Status: Informational

Expires: 25 April 2024

Authors: M. McFadden, Ed.

internet policy advisors, ltd

A Taxonomy of Internet Consolidation

Abstract

This document contributes to the ongoing discussion surrounding Internet consolidation. At recent IETF meetings discussions about Internet consolidation revealed that different perspectives gave completely different views of what consolidation means. While we use the term consolidation to refer to the process of increasing control over Internet infrastructure and services by a small set of organizations, it is clear that that control is expressed through economic, network traffic and protocol concerns. As a contribution to the discussion surrounding consolidation, this document attempts to provide a taxonomy of Internet consolidation with the goal of adding clarity to a complex discussion.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 25 April 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction - Why a Taxonomy?](#)
- [2. Toward a Definition of Consolidation](#)
- [3. Background to Consolidation Issues](#)
- [4. Economic Consolidation](#)
 - [4.1. 4.1. Economic Revenue Consolidation](#)
 - [4.2. 4.2. Economic Flow Consolidation](#)
- [5. Traffic and Infrastructure Consolidation](#)
- [6. Architectural Consolidation](#)
 - [6.1. The Rise of Intermediaries](#)
 - [6.2. Vertical Architectural Consolidation](#)
 - [6.3. Standards Development](#)
- [7. Service and Application Consolidation](#)
- [8. IANA Considerations](#)
- [9. Security Considerations](#)
- [10. Conclusions](#)
- [11. References](#)
 - [11.1. Informative References](#)
- [Acknowledgements](#)
- [Author's Address](#)

1. Introduction - Why a Taxonomy?

Internet consolidation has been under discussion for the last several years. The 2019 Internet Society's "Global Internet Report: Consolidation and the Internet Economy" highlighted issues on this topic and kicked started a series of discussions and publications around consolidation. The DINRG Workshop on Centralization took place in June of 2021 and was reported on at IETF 114[[DINRG](#)]. Furthermore, a draft for the Internet Architecture Board (IAB) discussed issues of economic and technical consolidation[[IABwrkshp](#)]. Despite community interest, the draft expired without further work or publication.

Several other contributions have focused on responses to consolidation. However, the report from the DINRG Workshop makes clear that there are different "categories of centralization." This draft does not attempt to propose responses to centralization, instead it attempts to build upon the Workshop's summary of "categories of centralization" with a goal of supporting future work and discussion.

2. Toward a Definition of Consolidation

There's a reason that concerns have been expressed about the "definition" of centralization. Centralization has seen many attempts to define "centralization." In a now expired Internet-Draft, one definition was "the process of increasing control over internet infrastructure and services by a small set of organizations." [\[Arrko1\]](#) A more recent Internet-Draft uses the following definition: "centralization" is the state of affairs where a single entity or a small group of them can observe, capture, control, or extract rent from the operation or use of an Internet function exclusively."[\[Nottingham\]](#)

This is not simply a recent trend. Dmytri Kleiner views the situation as primarily economic and notes that the dot com boom "was characterized by a rush to own infrastructure, to consolidate independent internet service providers and take control of the network." He describes the environment in the early 2000s as a sort of economic land rush where investors tried to replace the smaller service vendors with larger ones on every scale, from low level telecommunications infrastructure to high level services such as news aggregation, e-mail and video.[\[Kleiner\]](#)

Taking as a starting point the centralization of DNS services, there have also been studies of what centralization means in the context of key Internet services. For instance, one academic study found that about 12,000 name servers used by websites in the Alexa top 1 million used exactly the same third-party infrastructure. [\[DNSmeasure\]](#) This finding led to the proposal of a metric for measuring the market share of organizations that provide DNS resolvers.[\[Entropy\]](#) In this case, the authors implicitly provide definitions of "centralization" based on external metrics.

A completely different approach to the definition of consolidation can be seen by looking at research that focuses on the consolidation of data instead of economic measures.[\[DataConsolidation\]](#) In one study the authors state, "as users get used to using their services, the users' generated content and the data about their online behaviors are concentrated in such companies. This phenomenon, called "data consolidation", has become a serious problem, which makes the Internet society seek to decentralize the current Internet. The decentralized Internet aims to prevent the concentration of user data in a few giant companies."[\[ISOCconsolidation\]](#) This approach to centralization isn't an isolated research activity. Another study surveys the ongoing research activities on "data centralization" from the Internet Society, EU Horizon-2020, and Decentralized web.

The key problem is that "centralization" means different things in different contexts. Research takes on various topics including economic, network and data centralization. As a result the guidance from the DINRG workshop seems to have an important message for those who would like to see a single definition of centralization: there are "categories" of centralization, each with its own definition. [[DINRG](#)]

3. Background to Consolidation Issues

Internet consolidation is defined in one setting as "the process of increasing control over internet infrastructure and services by a small set of organizations." [[Nottingham](#)] Economy of scale is the driving force behind consolidation because markets naturally consolidate when economies of scale come into play.

The DINRG Workshop Report notes that:

- *the economy of scale enables one to generate the same service outcome with far lower production costs, and consume fewer resources for each instance of the service transaction.
- *a large user pool produces big data which helps improve service customization for each user, letting bigger companies gain an edge over smaller competitors.
- *centralized application developments reduce the number of platforms, hence substantially reduce the cost in development and maintenance, and circumvent interoperability issues. Consolidated development and operational efforts also help mitigate technical expertise shortages.
- *most of all, monopoly players can dictate to the market the terms of the service and the service price bought to the market, which causes longer term stagnation of the market and increased inefficiency within the market, which acts as a drag on further innovation.

The current consolidation and centralization of control and operation of Internet infrastructure and services was not an original design goal. In fact, RFC1958 says:

"This allows for uniform and relatively seamless operations in a competitive, multi-vendor, multi-provider public network."

and later,

"Heterogeneity is inevitable and must be supported by design." [[RFC1958](#)]

Much of the discussion surrounding consolidation focuses on Internet services, applications and data. However, contributions to the discussion of consolidation have also addressed economic, traffic and architectural consolidation. In recognition of this, a taxonomy with four main categories is proposed:

*Economic Consolidation

*Traffic and Infrastructure Consolidation

*Architectural Consolidation

*Service and Application Consolidation

4. Economic Consolidation

The Internet Society Report on Internet Consolidation suggests that the Internet's economy is defined as the economic activities that either support the Internet or are fundamentally dependent on the Internet's existence. [[ISOCconsolidation](#)]

As a result, economic consolidation on the Internet refers to the effects of market consolidation on competition and the economic power of a small set of companies that dominate economic activity in the Internet. There are two aspects to economic consolidation on the Internet.

Economic consolidation means that a small number of companies dominate the marketplace and hence, the revenues gathered from the use of the Internet.

Economic consolidation also means that a small number of companies control the flow of capital among enterprises that provide services on the Internet.

4.1. 4.1. Economic Revenue Consolidation

One of the two aspects of economic consolidation is the generation of revenue by a small number of enterprises. As an example, Amazon accounts for more than 45% of all online retail spending in the United States. Alibaba is estimated to have 60% of the electronic commerce market in China. Meta - including Facebook, Messenger, WhatsApp and Instagram - dominates social media and messaging holding four of the world's top six social media platforms.

In each of these cases, a very small number of companies operate extremely popular services, concentrating revenue generation into those companies. In fact, the popularity of these services is so great that value is created by adding other, complementary services onto the base they provide. The dominant services thus control the

foundation upon which other revenue streams are built.,xref
target="EmpiricalView"/>

The impact of revenue consolidation is enhanced by the presence of network effects in digital media and services. A network effect is market feature where the benefit to the user or consumer of a product or service increases as the number of other users increases. First-mover advantages play a role in defining the marketplace, but as a product or service gains traction and sees its market grow, the user base becomes the foundation for future growth.

An example of this would be the dominance of the portable document format, originally introduced by Adobe. As the number of users grew, it became the dominant format for the exchange of documents readable on a variety of devices. While the network effect does not imply immediate or inevitable market dominance, it does seem to be tightly correlated with market leaders in digital, Internet-based services.

4.2. 4.2. Economic Flow Consolidation

In addition to revenue generation, the very large application layer companies (Alphabet, Amazon, Tencent, Meta and Alibaba) control how money and capital moves through enterprises providing services on the Internet.

We have seen that embedded intermediaries that have substantial power can implement platforms that provide services to downstream consumers as well as upstream sources of content and applications. As the controlling intermediary for those services and applications, these large companies are also able to dictate how economic flows move between consumers and providers of applications and services.

Regulators and policy makers often are concerned about the enormous market power that these huge intermediaries have, but refrain from imposing controls or sanctions on the grounds that consumers get significant benefits when platform operators use upstream revenues to subsidize downstream services.

That control is often influenced by the fact that an enormous amount of the economic flow related to the Internet occurs largely out of public sight. Advertising is a critical part of the economic flow of the Internet. Advertisers want their message to reach the broadest community of potential consumers possible. As a result, advertising platforms need to be built to provide services to the greatest number of potential advertisers. The greater the reach of the advertising platform, the greater the value to the advertiser. The result, on the Internet, is a trend to collect the advertising tools in a consolidated group of global, dominant enterprises.

The end user sees free tools, services and information. However, the cost of those services is paid for, largely, by advertising revenue that makes up a vast proportion of the economic flow of the Internet.

5. Traffic and Infrastructure Consolidation

A significant majority of the Internet's traffic is delivered from very large content services including Google, Amazon and Facebook. These companies naturally attempt to provide the best possible service for their customers - including perceived speed of content delivery - these content services seek to establish connections directly with the companies providing access to the network. The result is a "flattening" of the Internet's traditional topology.

In fact, a recent study shows that these large services can reach more than 76% of the Internet without having to traverse traditional Tier 1 and Tier 2 ISPs. Besides bringing benefits of low latency and higher security to their users, these large-scale networks are also able to implement improvements and innovations in protocol design by having far greater control over the elements of the infrastructure being used to deliver services.

An empirical view of this consolidation in February of 2022 shows that the number of webpages that are hosted on these networks has increased from 2015 to 2020 at a rate exceeding 80%. In looking at data sources including TLD datasets and Alexa Top 1M datasets only a small number of content delivery networks host the vast majority of landing pages.

Centralization of this sort makes traffic filtering easier since forcing a content network to block specific content (or worse, blocking the content network entirely) would make a large amount of the content unavailable. As these networks begin to migrate other services to HTTP (for instance, DNS over HTTP), more than the web is affected by the impacts of filtering by centralized content services. In fact, blocking a content network entirely would block all the content of the network, not just the content that was the target of the filtering.

This happens at all layers. As an example at the application layer, in 2021 Google and Apple were forced to remove applications created by the Russian political opposition from their stores. The ability of a government to influence the content network means that centralization can lead to a reduction in the diversity of information or services on the Internet.

As the content networks grow in scale, the networks themselves grow to support the required network capacity. This sets up a feedback

look that drives market concentration toward the infrastructure provided by the content networks. As these networks grow larger, it becomes difficult for smaller networks and infrastructure providers to compete with the economies of scale from which the large networks benefit.

6. Architectural Consolidation

A third category of consolidation is the evolution of the Internet's architecture to meet contemporary use cases and requirements. Early descriptions of the Internet's architecture described heterogeneous endpoints connected by neutral transports. The end-to-end principle suggested that the transport of data between endpoints was provided without much intervention. [[RFC1958](#)]

Two developments have led to architectural consolidation: the emergence of intermediary services and the movement of transport related code to the application layer.

6.1. The Rise of Intermediaries

In the first case, technologies like CDNs are built into the network for the efficient delivery of content and services. The consumer is largely unaware that the service or application is being hosted by an organization other than the one they think they contacted. Instead, content delivery is pushed as close to the consumer as possible to ensure that the end-user experience is as optimal as possible.

The result is a series of security, economic and policy concerns associated with the small number of very large providers of these intermediary services. However, in this section we only want to consider the architectural issues specific to the use of intermediaries.

6.2. Vertical Architectural Consolidation

The second case is vertical architectural consolidation. This is where the companies that control the applications attempt to control all aspects of the communication. For instance, the provider of the browser may be the organization that the browser connects to. The advantage of this kind of architectural consolidation is that it allows the largest players to introduce technological innovation more quickly than if multiple layers of the stack required innovation in parallel.

With tools like DNS over HTTPS, we see applications taking control of the infrastructure of transport in addition to providing an application or service. Applications essentially provide their own

ecosystem (from centralized control of DNS services all the way to the end-user experience).

6.3. Standards Development

Others have rightly observed that, in the current environment, development of protocols and standards for the Internet is largely confined to a small number of participants from a small number of organizations. One trend is that the giant enterprises on the Internet also dominate the development of protocols.

Having a small number of organizations controlling the infrastructure of the Internet also means that innovative technologies can be implemented quickly and at large-scale. In 2022, a study in the ACM Transactions on Internet Technology found that Google accounting for 60% of all TLS 1.3 secured resources. Some other, large CDNs use TLS 1.3 almost exclusively. QUIC is also an example of a new technology that profits from consolidation. Large scale intermediaries can facilitate the deployment and adoption of new standards because the decision for the adoption is propagated across the infrastructure instead of through user adoption or feature updates.

7. Service and Application Consolidation

Services and applications are those tools that users see when they interact with the Internet. They take advantage of the infrastructure (and, access) parts of the Internet's ecosystem. According to the Internet Society's report on consolidation, "a small number of companies operating some of the Internet's most popular services dominate this market. Many of these companies act as multi-sided markets or platforms, meaning they offer a base upon which other applications, processes or technologies can be developed."

By itself, Google holds 90% of the global search market, the number one mobile operating system (Android), the top-user-generated video platform and has more than 1.5 billion active users of its Gmail email service. Google also has a map service, a public DNS resolver service, a cloud service and a document store.

While twenty years ago, an application would simply rely on the underlying operating system to provide its communications and transport services, now applications and services do this for themselves. This is a case of the intermediary or platform providing the application integrating all the necessary components for providing a service on the Internet.

8. IANA Considerations

This memo includes no request to IANA.

9. 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 6 above) have inserted new interfaces which must be reflected in security considerations for new protocols.

10. 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.

11. References

11.1. Informative References

- [Arrko1] Arrko, J., "Considerations on Internet Consolidation and the Internet Architecture", 2019, <<https://datatracker.ietf.org/doc/html/draft-arkko-arch-infrastructure-centralisation-00>>.
- [DataConsolidation] Kwon, T., Song, J., Jung, H., Chun, S., Lee, H., Kang, M., Park, M., and E. Cho, "How To Centralize the Internet: A Focus on Data Consolidation and User Privacy", 2022.
- [DINRG] Internet Engineering Research Group: DINRG, "Draft Report of DINRG Workshop on Centralization in the Internet", 2021, <<https://datatracker.ietf.org/meeting/114/materials/slides-114-dinrg-draft-report-of-dinrg-workshop-on-centralization-in-the-internet-01.pdf>>.
- [DNSmeasure] Zembruzki, L., Jacobs, A. S., Granville, G. S., and G.C.M. Moura, "Measuring Centralization of the DNS Infrastructure in the Wild", 2022.
- [EmpiricalView] Viet Doan, T., Rijswijk-Deij, R., Hohlfeld, O., and V. Bajpai, "An Empirical View on Consolidation of the Web", 2022, <<https://vaibhavbajpai.com/documents/papers/proceedings/web-consolidation-toit-2022.pdf>>.

[Entropy]

Bates, S., Bowers, J., Greenstein, S., Weinstock, J., Xu, Y., and J. Zittrain, "Evidence of Decreasing Internet Entropy;: The Lack of Redundancy in DNS Resolution by Major Websites", 2022.

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

[ISOCconsolidation] Internet Society, "Considerations on Internet Consolidation and the Internet Architecture [draft-arkko-iab-internet-consolidation-02] (Expired)", 2019, <<https://future.internetsociety.org/2019/consolidation-in-the-internet-economy>>.

[ISOCpaper] The Internet Society, "Consolidation In the Internet Economy", 2019, <<https://future.internetsociety.org/2019/consolidation-in-the-internet-economy>>.

[JCP] Journal of Cyber Policy, "Consolidation of the Internet", 2020, <<https://www.tandfonline.com/toc/rcyb20/5/1>>.

[Kleiner] Kleiner, D., "The Telekommunist Manifesto", 2022.

[Nottingham] Nottingham, M., "Internet Centralization: What Can Standards Do?", 2023, <<https://www.ietf.org/archive/id/draft-nottingham-avoiding-internet-centralization-14.html>>.

[RFC1958] Carpenter, B., "Architectural Principles of the Internet", RFC 1958, DOI 10.17487/RFC1958, June 1996, <<https://www.rfc-editor.org/info/rfc1958>>.

[RFC8890] Nottingham, M., "The Internet is for End Users", RFC 8990, DOI 10.17487/RFC8990, August 2020, <<https://www.rfc-editor.org/info/rfc8890>>.

Acknowledgements

Many thanks to all who discussed this with us in DINRG in 2021, 2022, and 2023.

Author's Address

Mark McFadden (editor)
internet policy advisors, ltd
6 Bridge Street
Chepstow

NP16 5EY
United Kingdom

Phone: [+1 608 504 7776](tel:+16085047776)

Email: mark@internetpolicyadvisors.com