

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 13, 2019

J. Arkko
Ericsson
B. Trammell
ETH Zurich
M. Nottingham
Fastly
C. Huitema
Private Octopus Inc.
M. Thomson
Mozilla
J. Tantsura
Nuage Networks
N. ten Oever
University of Amsterdam
March 12, 2019

**Considerations on Internet Consolidation and the Internet Architecture
draft-arkko-iab-internet-consolidation-01**

Abstract

Many of us have held a vision of the Internet as the ultimate distributed platform that allows communication, the provision of services, and competition from any corner of the world. But as the Internet has matured, it seems to also feed the creation of large, centralised entities in many areas. This phenomenon could be looked at from many different angles, but this memo considers the topic from the perspective of how available technology and Internet architecture drives different market directions.

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 <http://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 September 13, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

1. Introduction

Many of us have held a vision of the Internet as the ultimate distributed platform that allows communication, the provision of services, and competition from any corner of the world. But as the Internet has matured, it seems to also feed the creation of large, centralised entities in many areas.

We use the term Internet consolidation to refer to the process of the increasing control over Internet infrastructure and services by a small set of organizations. Such concentration has an obvious effect on traffic flows or whose services and systems are in daily use by large population of Internet users, but it can also create secondary effects, where the ability to collect information or to affect something is concentrated in the small set of organizations.

Consolidation may also affect technology choices and the evolution of the Internet architecture, e.g., large organizations or organizations providing important technology components may have a significant impact on what technology is deployed for large numbers of users or by other organizations.

Our first question is whether Internet is indeed consolidating? It certainly appears so, though more quantitative research on this topic would be welcome. It is also possible that there's only a perception of consolidation, as market forces have caused business changes in new areas of business. Arguably, today's consolidation areas seem to be more in the application space than further down in the stack or in the operating systems, as was the case some years ago. The second question is if we're seeing consolidation simply moving to new areas.

This phenomenon could be looked at from many different angles, but this memo considers the topic from the perspective of how available

technology and Internet architecture drives different market directions. Our third question is if the Internet technology has influenced the consolidation trends in some manner? And conversely, the fourth question asks how is Internet consolidation influencing the development of the Internet infrastructure and architecture?

The engineering remit at the IETF is to focus on technology, but of course we also want to understand the implications and externalities of the technical arrangements we design. Technology affects economics and vice versa. The Internet technology community continues to make decisions that have ramifications on Internet systems, just as we are subject to forces that affect them.

As technologists, our fourth question is whether there are changes in technology that would help reduce those large-player advantages that are technically-driven.

This memo reviews areas where consolidation may be occurring in the Internet, and discusses the potential reasons for this. The memo starts by review other work in this area in [Section 2](#). [Section 3](#) discusses consolidation and the reasons behind the creation of larger entities, and [Section 4](#) looks at some actions that might alleviate the situation.

If you are interested on this or other architecture-related topics, please subscribe to the IAB architecture-discuss mailing list as one forum for discussion. Similarly, the Internet Society has chosen consolidation as a focus topic for their year 2019 activities. Their report is in [\[ISOC\]](#).

[2.](#) Other Work

One of the causes for the current consolidation of the Internet infrastructure can be traced back to some of the assumptions that were made during the commercialization of the Internet in the early 1990s [\[Abbate\]](#), even though [\[RFC1192\]](#) already describes some potential issues that could arise. Overall it was expected the combination of commercialization, together with the technical and architectural characteristics of the Internet, such as its modularity and layering principles, would lead to perfect markets, free competition and decentralized structures [\[LitanRivlin\]](#).

But as we know now, this did not happen entirely as expected. Some even argue that 'market concentrations, control and power struggles are categories to adequately describe the fundamental dynamics of the commercial Internet' [\[DolataSchrape\]](#). While the privatization was supposed to lead to competition and innovation [\[Cowheyetal\]](#)

[[VanSchewick](#)], some argue that it actually led to the emergence of Internet oligopolies [[Mansell](#)] [[Smyrnaiois](#)].

Current scientific economic thinking harbors two different schools vis a vis efficient markets and monopolies. The school based on Adam Smith argues that unfettered markets tend to concentration of wealth and income, whereas liberal economist believe in efficient markets that stimulate competition.

On the other hand, according to Joseph Stiglitz 'many sectors - telecoms, cable TV, digital branches from social media to Internet search, health insurance, pharmaceuticals, agro-business, and many more - cannot be understood through the lens of competition' [[Stiglitz](#)]. The considerations of technologists and policy makers at the time of the commercialization and privatization of the Internet infrastructure might have been based on a belief in efficient markets, whereas we are now finding out this might not always be how markets function.

It cannot be denied that 'market actors have contributed immensely to the evolution of the Internet in terms of investment, products, services, and infrastructure, and the government's light-handed approach to regulation has given producers and consumers substantial freedom to innovate and to self-regulate with respect to many issues affecting the Internet community in ways that have produced substantial social benefits' [[Frischmann](#)]. But the current consolidation in ownership of and control over the Internet infrastructure was not foreseen [[Clark](#)], and arguably the loss of decentralized control goes against its design objectives. For instance, [[RFC1958](#)] says:

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

and

Heterogeneity is inevitable and must be supported by design.

And [[RFC3935](#)] says:

We embrace technical concepts such as decentralized control, edge-user empowerment and sharing of resources, because those concepts resonate with the core values of the IETF community.

3. Factors Driving Consolidation

Consolidation is driven by economic factors relating to scale and ability to easily reach a large market of users over the Internet. This kind of setting tends to enable winners to take large market shares, whether those winners came about through the model that liberal economists believe in or the model that Adam Smith believes in.

The most visible aspects of this involve well-recognized Internet services. Internet Society's report summarised the market position of popular Internet service brands as follows [[ISOC](#)]:

- o Facebook and Google have been estimated to account for 84% of global digital advertising investment (excluding China).
- o Amazon is expected to account for 49.1% of all online retail spending in the US. Similarly, Alibaba is estimated to have close to 60% of the e-commerce market in China.
- o Google alone holds 90% of the global search market, over 60% of web browsers, and has the number 1 (by far) mobile operating system (Android).
- o Facebook - incorporating Facebook Messenger, WhatsApp, and Instagram - holds 4 of the world's top 6 social media platforms.

But it is important to recognize that the Internet is a complex ecosystem. There are many underlying services whose diversity, or lack thereof, are as important as that of, say, consumer-visible social networks. For instance, the diversity of cloud services, operating systems, browser engines is as important as that as of application stores or the browsers themselves.

Of course, the Internet allows plenty of choice both in these and other areas. Too many or too few choices create different kinds of problems.

It would be useful to break these general factors and observations down a bit further. In particular, it is useful to distinguish market or economic factors from technical factors.

3.1. Economics

Scaling benefits are natural for many types of businesses. And many Internet-based businesses can potentially serve a very large customer base, as the cost of replicating and delivering their service to new customers or areas is small.

However, typically the network effect has an even more pronounced impact. Each additional user adds to the value of the network for all users in a network. In some applications, such as the open web, this value grows for everyone, as the web is a globally connected, interoperable service for anyone with a browser can use.

There is an important distinction between different applications of the network effect, however. Consider email as another example; anyone with an account at any email server can use it globally. However, here we have seen much more consolidation into few large email providers, both due to innovative, high-quality services but also because running email services by small entities is becoming difficult; among other things due to spam prevention practices that tend to recognize well only the largest entities.

In some other applications, such as social media, the services have a more closed nature. The value of being a customer of one social media service depends highly on how many other customers that particular service has. Hence, the larger the service, the more valuable it is. And the bigger the value difference to the customers, the less practical choice they have in selecting a service.

In some cases, these developments also allow asymmetric relationships to form, with the customers having less ability to affect the service than they would perhaps wish.

[3.2.](#) Interdependencies

Entities with a large role in a market place tend to have inertia, of course through having many customers, but also due to their interconnectedness with the rest of the ecosystem. These interconnections can range from business relationships to peering connections, linking, and the use of various infrastructure services from an entity as building blocks in applications.

These interconnections make it difficult for a market to move away from a large entity. Indeed, even for an individual it can be a challenge to not use the most commonly used Internet services [[Gizmodo](#)].

Optimistic technologists ("digital libertarians") tend to have a beliefs that states have limited ability to regulate the Internet: "The Net interprets censorship as damage and routes around it" [[Gilmore](#)]. However, as argued by [[Boyle](#)] states may have multiple ways to influence and monitor the Internet. One of the issues related to consolidation is that it tends to be easier to exert control of few large entities, than a large set of small, distributed

players. This concern is particularly acute around intellectual property rights or surveillance capabilities, particularly when extra-territorial requirements are placed on the large entities. These entities cannot avoid conforming to regulation and laws in any of the locations they have presence in.

As a result, there's an added angle of interconnectedness with governments. At the same time, this of course also provides an avenue for control of market forces, e.g., in the form of competition law.

3.3. Data- and Capital-intensive Services

The scaling advantages are only getting larger with the advent of AI- and machine learning -based technologies.

The more users a service has, the more data is available for training machine learning models, and the better the service becomes, bringing again more users. This feedback loop and the general capital-intensive nature of the technology (data and processing at scale) makes it likely that the largest companies are ahead in the use of these technologies.

One could also take the pessimistic view that many of the ongoing disputes in standards organizations relate to which market actors will ultimately be able to collect the more data from private persons and how. The question isn't as much about the protection of the user's privacy but rather whether one or another industry (or country) will be able to benefit from access to data.

3.4. Permissionless Innovation

The email vs. social media example also highlights the interesting roles of interoperability and the "permissionless innovation" principle - the idea that a network can be simple but still powerful enough that essentially any application could be built on top of it without needing any special support from anyone else. Permissionless innovation has brought us all the innovative applications that we enjoy today, on top of a highly interoperable underlying network, along with advances in video coding and other techniques used by applications.

Paradoxically, if the underlying network is sufficiently powerful, the applications on top can evolve without similar pressures for interoperability, leading to the closed but highly valuable services discussed above. We call this the Permissionless Completeness Problem.

3.5. Fundamentals of Communication

There are also fundamental issues. For instance, speed of light; low-latency services can fundamentally only be provided through globally distributed data centers. These are often provided built by large organizations, although collaborative and data center or cloud computing service approaches also exist.

A similar issue has arisen in recent years around large-scale denial-of-service attacks, and how various entities can deal with them. While the largest attacks affect all players (see, for instance, the Dyn attacks in October 2016), it is also true that large cloud- and content delivery providers can better deal with such attacks due to their scale. This is one reason that attracts many network services to such providers.

3.6. Technology Factors

One of the key questions is whether we are seeing developments that are driven by economic factors or whether fundamental reasons or lack available technology drives particular models. For instance, centralised solutions might be desirable due to business incentives, or they might be necessary because there is no distributed, collaborative solution.

For instance, some technical issues have historically not been easy to solve, such as e-mail spam, which has led to reliance on non-technical solutions. Today, it is becoming increasingly difficult to run your own mail services, essentially forcing many organizations and individuals to employ larger providers. The issues relate directly to size of entities; no one can afford to disconnect from the largest providers. But as a small entity, there is little leverage to convince peer entities or various supporting white/blacklist entities to deal with you properly.

Many Internet services are based on gathering data about users, and using that data for, for instance, targeted advertisements. More data from more users makes it possible to run a service more accurately or with better results; here again scale brings advantages.

Another trend is that more and more content is becoming available locally, from a content delivery or provider function directly on your own ISP's network. We predict that eventually most content will be delivered this way, reducing the role that global IP connections across the Internet play. By some metrics this has already happened; what practical - positive or negative - impacts might this have on the Internet technology?

There are also security tradeoffs. Large entities are generally better equipped to move to more recent and more secure technology. For instance, the Domain Name System (DNS) shows signs of ageing but due to the legacy of deployed systems, has changed very slowly. Newer technology developed at the IETF enables DNS queries to be performed confidentially, but its deployment is happening mostly in browsers that use global DNS resolver services, such as Cloudflare's 1.1.1.1 or Google's 8.8.8.8. This results in faster evolution and better security for end users.

However, if one steps back and considers the overall security effects of these developments, the resulting effects can be different. While the security of the actual protocol exchanges improves with the introduction of this new technology, at the same time this implies a move from using a worldwide distributed set of DNS resolvers into, again, more centralised global resolvers. While these resolvers are very well maintained (and a great service), they are potentially high-value targets for pervasive monitoring and Denial-of-Service (DoS) attacks. In 2016, for example, DoS attacks were launched against Dyn, one of the largest DNS providers, leading to some outages.

4. Call to Action

Are there assumptions about the Internet architecture that no longer hold in a world where larger, more centralised entities provide big parts of the Internet service? If the world changes, the Internet and its technology/architecture may have to match those changes.

It appears that level the playing field for new entrants or small players brings potential benefits. Are there technical solutions that are missing today?

Of course, it may well be that technology improvements are hard to come by. Nevertheless, recognizing the risks of consolidation in both current and proposed future technologies is the first step in proactively avoiding those risks where possible.

Assuming that one does not wish for regulation, technologies that support distributed architectures, open source implementations of currently centralised network functions, or help increase user's control can be beneficial. Federation, for example, would help enable distributed services in situations where smaller entities would like to collaborate.

Similarly, in an asymmetric power balance between users and services, tools that enable the user to control what information is provided to a particular service can be very helpful. Some such tools exist, for

instance, in the privacy and tracking-prevention modes of popular browsers but why are these modes not the default, and could we develop them further?

It is also surprising that in the age of software-defined everything, we can program almost anything else except the globally provided, packaged services. Opening up interfaces would allow the building of additional, innovative services, and better match with users' needs.

Silver bullets are rare, of course. Internet service markets sometimes fragment rather than cooperate through federation. And the asymmetric power balances are easiest changed with data that is in your control, but it is much harder to change when someone else holds it. Nevertheless, the exploration of solutions to ensure the Internet is kept open for new innovations and in the control of users is very important.

- o What IETF topics that should be pursued to address some of the issues around consolidation? Some of the topics for consideration are discussed in [Section 4.1](#) and [Section 4.2](#).
- o Are there open source efforts that should be pursued or kept being pursued to ensure that a diversity of operators and networks can use a particular technology? This is further discussed in [\[open-source\]](#).
- o What measurements relating to the developments centralization or consolidation should be pursued? And what other research, such as distributed Internet architectures, should be driven forward? Some potential topics are discussed in [Section 4.4](#).

[4.1](#). Open Interfaces

Standards and open source efforts continue to build many open interfaces and APIs that allow systems interoperability and tailoring of services. In some cases, however, the availability of open interfaces definitions and software has not lead to the realization of actual interfaces in this open manner. For instance, different instant messaging systems have had a technical ability to be interoperable with other systems, just like e-mail is interoperable across systems, but have chosen to be disconnected.

Work in determining what open interfaces can provide benefits to users as well be successfully deployed in the Internet ecosystem would of course be useful.

4.2. Specific Standardization Choices

Sometimes the issue is not the availability of interfaces as such, but rather fundamental architectural choices with regards to how Internet systems should be built.

Often this relates to how centralized or distributed deployments are targeted. And even if a distributed, broad deployment model is targeted, the expectations may not match reality when economies of scale, DDoS resilience, market consolidation, or other factors have come into play. These factors can result in the deployed reality being highly concentrated.

This is a serious issue for the Internet, as concentrated, centralized deployment models present risks to user choice, privacy, and future protocol evolution. On occasion, the differences to expectations can be immediate, but can also occur much later.

Some examples of these issues include current work in DNS where we may see concentration through the use of globally available common resolver services, which evolve rapidly and can offer better security. But the concentration of these queries into few services creates new security and privacy concerns.

Another example is email, which started out as many providers running in a largely uncoordinated fashion, but which has since then seen significant market consolidation and a need for coordination to defend against attacks such as spam. The coordination and centralized defense mechanisms scale better for large entities, which has fueled additional consolidation.

Awareness of these issues while working on standards would be useful, so that the issues can be taken into account and appropriately mitigated. To begin with, those of us interested in the broader questions about Internet development and Internet researchers are needed to identify some of the effects that new technology developments may pose.

4.3. Open Source

Many key Internet infrastructure services (e.g., DNS servers), end-user applications (e.g., browsers) and technology components (e.g., operating systems or protocol implementations) are commonly implemented using open source solutions.

This is often true even if there's a large entity that is in charge of large fraction of development and deployment for a particular technology. This is a good thing, as it ensures that others have a

fair chance of changing the technology in question, should the large entity drive their use of the technology in some direction that does not benefit the users. For instance, users and other organizations have traditionally been able to either run their own browser versions or provided extensions that suit their needs better than the default system.

As a result, continuing to have this ecosystem is important safety valve and competition opportunity.

4.4. Research Challenges

There are number of different research directions for which further work would be useful.

The first is about measurements; is the Internet indeed consolidating, and if so, by how much, and in what aspects? Also, where are Internet's traffic flows concentrating, and how is this changing over time?

There are also questions about the trends and their relationships to technology: Has Internet technology influenced the consolidation trends in some manner? And conversely, how Internet consolidation influences the development of the Internet infrastructure and architecture?

Finally, research on topics that would likely yield results that increase the diversity and de-centralization in the Internet would obviously be welcome. This can include any de-centralized technologies, but in particular distributed Internet architecture is an Interesting topic.

5. Contributors

Much of the text in this memo is from a blog article written by Jari Arkko, Mark Nottingham, Christian Huitema, Martin Thomson, and Brian Trammell for the Internet Architecture Board (IAB), and from a blog article written by Jari Arkko and Brian Trammell APNIC and RIPE. Some parts of the text have also come from a future workshop description developed in the IAB, primarily by Christian Huitema and Ted Hardie.

6. Acknowledgements

The authors would like to thank IAB members, Geoff Huston, Amelia Andersdotter, Gonzalo Camarillo, Mirjam Kuehne, Robert Mitchell, Olaf Kolkman, and many others for interesting discussions in this problem space.

7. Informative References

- [Abbate] Abbate, J., "Inventing the Internet", MIT Press , 2000, <<https://mitpress.mit.edu/books/inventing-internet>>.
- [Boyle] "Foucault in Cyberspace: Surveillance, Sovereignty, and Hardwired Censors", Duke Law School, https://scholarship.law.duke.edu/faculty_scholarship/619/ , 1997.
- [Clark] Clark, D., "The Design Philosophy of the DARPA Internet Protocols", In Symposium Proceedings on Communications Architectures and Protocols, 106-114. SIGCOMM '88. New York, NY, USA, ACM <https://doi.org/10.1145/52324.52336> , 1988.
- [Cowheyetal]
Cowhey, P., Aronson, J., and J. Richards, "Shaping the Architecture of the US Information and Communication Technology Architecture: A Political Economic Analysis", Review of Policy Research 26 (1-2), pp. 105-125. , 2009.
- [DolataSchrape]
Dolata, U. and J. Schrape, "Collectivity and Power on the Internet: A Sociological Perspective", Springer International Publishing. Page 85. , 2018, <www.springer.com/de/book/9783319784137>.
- [Frischmann]
Frischmann, B., "Privatization and Commercialization of the Internet Infrastructure", Colum. Sci. & Tech. L. Rev. 2, pp. 1-25. , 2001.
- [Gilmore] "First Nation in Cyberspace", TIME International (see <http://kirste.userpage.fu-berlin.de/outerspace/internet-article.html>) , December 1993.
- [Gizmodo] "I Cut The 'Big Five' Tech Giants From My Life. It Was Hell", Gizmodo, <https://www.gizmodo.com.au/2019/02/i-cut-the-big-five-tech-giants-from-my-life-it-was-hell/> , February 2019.
- [ISOC] "Consolidation in the Internet economy", Internet Society, <https://future.internetsociety.org/2019/> , 2019.

[LitanRivlin]

Litan, R. and A. Rivlin, "Projecting the Economic Impact of the Internet", American Economic Review 91 (2), pp. 313-317 , 2001.

[Mansell] Mansell, R. and M. Javary, "Emerging Internet Oligopolies: A Political Economy Analysis", In Miller, Edythe S. and Samuels, Warren J., (eds.) An Institutionalist Approach to Public Utilities Regulation. Michigan State University Press, East Lansing, Michigan, pp. 162-201. ISBN 9780870136245 , 2002, <<https://core.ac.uk/download/pdf/204420.pdf>>.

[RFC1192] Kahin, B., "Commercialization of the Internet summary report", [RFC 1192](#), DOI 10.17487/RFC1192, November 1990, <<https://www.rfc-editor.org/info/rfc1192>>.

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

[RFC3935] Alvestrand, H., "A Mission Statement for the IETF", [BCP 95](#), [RFC 3935](#), DOI 10.17487/RFC3935, October 2004, <<https://www.rfc-editor.org/info/rfc3935>>.

[Smyrnaios]

Smyrnaios, N., "Internet Oligopoly: The Corporate Takeover of Our Digital World", Emerald Publishing Ltd. , 2018.

[Stiglitz]

Stiglitz, J., "Joseph Stiglitz: Are markets efficient, or do they tend towards monopoly? The verdict is in", 2016, <<https://www.weforum.org/agenda/2016/05/joseph-stiglitz-are-markets-efficient-or-do-they-tend-towards-monopoly-the-verdict-is-in/>>.

[VanSchewick]

van Schewick, B., "Internet Architecture and Innovation", MIT Press , 2012.

Authors' Addresses

Jari Arkko
Ericsson

Email: jari.arkko@piuha.net

Brian Trammell
ETH Zurich

Email: ietf@trammell.ch

Mark Nottingham
Fastly

Email: mnot@mnot.net

Christian Huitema
Private Octopus Inc.

Email: huitema@huitema.net

Martin Thomson
Mozilla

Email: martin.thomson@gmail.com

Jeff Tantsura
Nuage Networks

Email: jefftant.ietf@gmail.com

Niels ten Oever
University of Amsterdam

Email: mail@nielstenoever.net

