Internet Centralization: What Can Standards Do?

Abstract

Despite the Internet being designed and operated as a decentralized network-of-networks, forces often (and increasingly) encourage consolidation of power over its functions into few hands.

This document discusses centralization in Internet protocols and relates it to consolidation of power, explains why both are undesirable, identifies forces that contribute to them, catalogues limitations of common approaches to decentralization, and explores what Internet standards efforts can do.

About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-nottingham-avoiding-internet-centralization/.

Source for this draft and an issue tracker can be found at https://github.com/mnot/avoiding-internet-centralization.

Status of This Memo

Table of Contents

1. **Introduction**

   The Internet has succeeded in no small part because of its
   purposeful avoidance of any single controlling entity. While
   originating in a desire to prevent a single technical failure from

having a wide impact [BARAN], this stance has also enabled the Internet's rapid adoption and broad spread. The Internet can accommodate a spectrum of requirements and is now a global public good because permission is not required to connect to it, deploy an application on it, or use it for a particular purpose.

While maintaining these properties remains a widely shared goal, consistently achieving them has proven more difficult over time. Today, many successful services operate in a centralized fashion -- to the point where some proprietary applications have become so well-known that they are commonly mistaken for the Internet itself. Even when open protocols incorporate techniques intended to prevent consolidation, economic and social factors can drive users to prefer solutions built with or on top of supposedly decentralized technology.

These difficulties call into question what role architectural design -- in particular, that overseen by open standards bodies such as the IETF -- should play in preventing, mitigating, and controlling consolidation of power on the Internet. This document discusses aspects of centralization that relate to Internet standards efforts, and argues that while the IETF may not be able to prevent these outcomes, there are still meaningful steps it can take to help counteract them.

Section 2 defines centralization and consolidation, explains why and when they are undesirable, and surveys how they occur on the Internet. Section 3 explores decentralization and highlights some relevant techniques, along with their limitations. Finally, Section 4 considers the role that Internet standards play in avoiding centralization and mitigating its effects.

The primary audience for this document is the engineers who design and standardize Internet protocols. However, designers of proprietary protocols and applications can benefit from considering these issues, especially if they intend their work to be considered for eventual standardization. Likewise, policymakers can use this document to help identify and evaluate proposed remedies for inappropriately centralized protocols and applications.

2. Centralization and Consolidation

This document distinguishes "consolidation" from "centralization" to separate effects from (some of) their causes.

"Consolidation" is 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.

Here, "entity" could be a single person, a corporation, or a government. It does not include an organization that operates in a manner that effectively mitigates consolidation (see, e.g., Section 3.1.2).

"Internet function" is used broadly in this document. It might be an enabling protocol already defined by standards, such as IP [RFC791], BGP [RFC4271], TCP [RFC793], or HTTP [HTTP]. It might also be a proposal for a new enabling protocol, or an extension to an existing one.

Furthermore, Internet functions are not limited to standards-defined protocols. User-visible applications built on top of standard protocols are also vulnerable to consolidation -- for example, social networking, file sharing, financial services, and news dissemination. Likewise, networking equipment, hardware, operating systems, and software act as enabling technologies that can exhibit consolidation. The supply of Internet connectivity to end users in a particular area or situation can also be subject to consolidation, as can the supply of transit between networks (so called "Tier 1" networks).

"Centralization" is the source of consolidation that this document focuses upon; it measures the direct or potential contribution of a function's technical design to consolidation. For example, many consider the social networking market to be highly consolidated around a few providers who have used highly centralized architectures (see Section 2.2.1) to reinforce their control.

Centralization is not a binary condition; a function's design might contribute to or be vulnerable to consolidation in multiple ways and in various degrees. Even when decentralization techniques are purposefully used to avoid centralization in a particular aspect of a function, it often appears in other places -- for example, in its governance, implementation, deployment, or in ancillary functions. In summary, "decentralized technology alone does not guarantee decentralized outcomes." [SCHNEIDERa] Therefore, we need a more nuanced way to asess it.

2.1. Assessing Centralization

By default, Internet protocol designers will avoid an obviously centralized design because the Internet's very nature is incompatible with it. As a "large, heterogeneous collection of interconnected systems" [BCP95] the Internet is often characterised as a "network of networks". These networks relate as peers who agree to facilitate communication, rather than having a relationship of subservience to others' requirements or coercion by them. This focus on independence of action carries through the way the network is

architected -- for example, in the concept of an "autonomous system".

However, as discussed below in Section 2.2.2, not all centralization is avoidable, and in some cases, it is even desirable. With that in mind, centralization on the Internet is most concerning when it is not broadly held to be necessary, when it has no checks, balances, or other mechanisms of accountability, when it selects "favorites" which are difficult (or impossible) to displace, and when it threatens to diminish the success factors that enable the Internet to thrive -- scalability to meet the demands of new users, adaptability to encompass new applications, flexibility to enable deployment of new technologies, and resilience to shocks and changes [KENDE].

Most often, unacceptable centralization is indicated when a proposal has one or more of the following damaging effects (or the potential for them):

  *Power Imbalance: When a third party has unavoidable access to
   communications, the informational and positional advantages
   gained allow observation of behavior (the "panopticon effect")
   and shaping or even denial of behavior (the "chokepoint effect")
   [JUDGE] -- capabilities that those parties (or the states that
   have authority over them) can use for coercive ends [FARRELL] or
   even to disrupt society itself. Just as good governance of states
   requires separation of powers [MADISON], so too does good
   governance of the Internet require that power not be concentrated
   in one place without appropriate checks and balances.

  *Limits on Innovation: Consolidation can preclude the possibility
   of "permissionless innovation" -- the ability to deploy new,
   unforeseen applications without requiring coordination with
   parties other than those you are communicating with.

  *Constraints on Competition: The Internet and its users benefit
   from robust competition when applications and services are
   available from many providers -- especially when those users can
   build their own applications and services based upon
   interoperable standards. When a consolidated service or platform
   must be used because no substitutes are suitable, it effectively
   becomes an essential facility, which encourages abuse of power.

  *Reduced Availability: Availability of the Internet (and
   applications and services built upon it) improves when there are
   many ways to obtain access. While service availability can
   benefit from the focused attention of a large consolidated
   provider, that provider's failure can have a disproportionate
   impact on availability.

Monoculture: The scale available to a consolidated provider can
  magnify minor flaws in features to a degree that can have broad
  consequences. For example, a single codebase for routers elevates
  the impact of a bug or vulnerability; a single recommendation
  algorithm for content can have severe social impact. Diversity in
  functions' implementation leads to a more robust outcome when
  viewed systemically, because "progress is the outcome of a trial-
  and-error evolutionary process of many agents interacting
  freely." [ALIGIA]

  Self-Reinforcement: As widely noted (see, e.g., [ABRAHAMSON]), a
  consolidated provider's access to data allows it the opportunity
  to make improvements to its offerings, while denying such access
  to others.

However, these are only indicators, and each needs to be evaluated
carefully on a case-by-case basis.

For example, it is important to distinguish centralization from
anticompetitive concerns (also known as "antitrust"). While there
are many interactions between these concepts and making the Internet
more competitive may be a motivation for avoiding consolidation,
only courts (and sometimes, regulators) have the authority to define
a relevant market and determine that behavior is anti-competitive.
Furthermore, what might be considered undesirable consolidation by
the technical community might not attract competition regulation.
Conversely, what might attract competition regulation might not be
of great concern to the technical community if other mitigations are
felt to be adequate.

Likewise, while centralization interacts with availability, they are
distinct and any relationship between them cannot be assumed without
careful analysis of where and how centralization occurs. Centralized
systems might be more available because of factors like the
resources available to them, but also have greater impact when they
encounter a fault; decentralized systems might be more resilient in
the face of local failures, but less able to react to systemic
issues. A failure because of a cut cable, power outage, or failed
server is qualitatively different from the issues encountered when a
core Internet function has a gatekeeper.

For example, a large variety of Web sites might depend upon a cloud
hosting provider or content delivery network; if it were to become
unavailable (whether for technical or other reasons), many people's
experience of the Internet might be disrupted. Likewise, a mobile
Internet access provider might have an outage that affects hundreds,
thousands, or more of its users. In both cases, centralization is
not indicated by the loss of availability or its scale, but it well
might be if the parties relying on the function don't have

reasonable options to switch to if they are unhappy with the
availability of the service provided, or if friction against
switching to an alternative is too great.

2.2. How Centralization Occurs

A function's design can exhibit and encourage centralization in a
variety of ways. The subsections below describe different
contributors to and expressions of centralization in Internet
functions. Note that this is not a taxonomy, in that it is not
complete and there may be overlap.

2.2.1. Proprietary Centralization

Creating of a protocol or application with a fixed role for a
specific party is the most obvious form of centralization. Many
messaging, videoconferencing, chat, social networking, and similar
applications currently operate in this fashion.

Because they allow control by a single entity, proprietary protocols
are often considered simpler to design, more amenable to evolution,
and more likely to meet user needs [MOXIE], compared to
decentralized alternatives. However, their centralization is
absolute -- if the function has no alternative providers, or
switching to those providers is too difficult, its users are "locked
in."

Proprietary protocols and applications are not considered as being
part of the Internet per se; instead, they are more properly
characterized as being built on top of the Internet. The Internet
architecture and associated standards do not control them, beyond
the constraints that the underlying protocols (e.g., TCP, IP, HTTP)
impose.

2.2.2. Beneficial Centralization

Some protocols and applications have goals that require
centralization, because they rely on it to deliver a particular
benefit.

Often, this is due to technical necessity. For example, a single,
globally coordinated "source of truth" is by nature centralized --
such as in the Domain Name System (DNS), which allows human-friendly
naming to be converted into network addresses in a globally
consistent fashion.

Or, consider IP addresses allocation. Internet routing requires
addresses to be allocated uniquely, but if a single government or
company were to capture the addressing function, the entire Internet
would be at risk of abuse by that entity. The same benefits and

risks can be seen in the Web's trust model, thanks to the Certificate Authority's role in communication between clients and servers

Protocols that need to solve the "rendezvous problem" to coordinate communication between two parties who are not in direct contact also exhibit beneficial centralization. For example, chat protocols need to coordinate communication between two parties that wish to talk; while the actual communication can be direct between them (so long as the protocol facilitates that), the endpoints' mutual discovery typically requires a third party at some point. From the perspective of those two users, the rendezvous function is centralized.

Even when not strictly necessary, centralization can be deployed to beneficial ends. [SCHNEIDERb] notes that "centralized structures can have virtues, such as enabling publics to focus their limited attention for oversight, or forming a power bloc capable of challenging less-accountable blocs that might emerge. Centralized structures that have earned widespread respect in recent centuries – including governments, corporations, and nonprofit organizations – have done so in no small part because of the intentional design that went into those structures."

For example, when traffic from many users is mixed in a way that can't be distinguished, censorship becomes more difficult. This "too big to block" phenomenon drives the design of many recent protocols (such as [ECH]), but they require a degree of centralization to meet their goals.

Likewise, when a function requires governance to realize common goals and protect minority interests, a "choke point" is naturally formed by the chosen governance mechanism, increasing centralization. For example, content moderation functions concentrate decision making to impose community values. Complex and risky functions like financial services (e.g., credit card processing) can be seen as beneficially centralized into relatively few, specialized organizations, where they can received the focused attention that they require.

When beneficial centralization is present, Internet protocols often attempt to mitigate the associated risks using measures such as federation (see Section 3.1.1) or multi-stakeholder governance (see Section 3.1.2). Protocols that successfully do so are often reused to avoid the considerable cost and risk of re-implementing those mitigations. For example, if a protocol requires a coordinated, global naming function, reusing the Domain Name System is usually preferable to establishing a new system.

Ultimately, deciding what is beneficial is a judgment call. Some
protocols cannot function without a centralized function; others
might be significantly enhanced for certain use cases if a function
is centralized, or might merely be more efficient. Such judgments
should be made in light of established architectural principles and
how benefits accrue to end users.

2.2.3. Concentration

Even when a function avoids or mitigates other forms of
centralization, it might become consolidated in practice when
external factors influence its deployment, so that few or even just
one entity provides the function. This document refers to this
phenomenon as "concentration." Concentration can be caused by
economic, legal, social, and even cognitive factors that encourage
use of a central function despite the absence of such a requirement
in the protocol itself.

Concentration is often associated with the network effects that are
so often seen on the Internet. While in theory every node on the
Internet is equal, in practice some nodes are much more connected
than others: for example, just a few sites drive much of the traffic
on the Web. While expected and observed in many kinds of networks,
these effects award asymmetric power to nodes that act as
intermediaries to communication. [BARABASI]

There may be legitimate qualitative reasons for some nodes being
favoured over others. However, when it happens because friction
against using an alternative prevents switching, benefits are
accrued to services rather than users. If choosing an alternate
provider requires a significant amount of time, resources,
expertise, coordination, loss of functionality, or effort,
centralization is indicated.

Conversely, a function based on a well-defined, open specification
designed to minimize switching costs might be considered to have
less centralization even when users continue to favor large
providers, because ease of switching creates implicit competitive
pressure upon them.

For example, social networking is an application that is currently
supplied by a few proprietary platforms despite standardization
efforts (see, e.g., [ACTIVITYSTREAMS]), because of the powerful
network effects associated. While there has been some competition in
social networking, the choices that their peers make often restricts
individual choices, because of the coordination required to move to
a new service.

See [ISOC] for a deeper exploration of concentration.

Concentration is difficult to avoid in protocol design, and
federated protocols are particularly vulnerable to it (see
[Section 3.1.1](#)).

2.2.4. Inherited Centralization

Most Internet protocols and applications depend on other, "lower-
layer" functions and their implementations. The features,
deployment, and operation of these dependencies can surface
centralization into functions and applications built "on top" of
them.

For example, the network between endpoints can introduce this kind
of centralization to application-layer protocols because it is
necessary for communication and therefore has power over it. A
network might block access to, slow down, or change the content of
various application protocols or specific services for financial,
political, operational, or criminal reasons, thereby creating a
disincentive (or even removing the ability) to use them. By
selectively hindering the use of some services but not others,
network interventions can be composed to aid concentration in those
other services -- intentionally or not.

Likewise, having only a single implementation of a protocol is a
form of inherited centralization, because applications that use it
are vulnerable to the control it has over their operation. Even Open
Source projects can exhibit this if there are factors that make
forking difficult (for example, the cost of maintaining that fork).

Inherited centralization surfaces when viable alternatives to these
dependencies are not available. It is often present when network
effects restrict choices, but can also be created by legal mandates
and incentives that restrict the options for a function (such as
Internet access), its provision, or the range of implementations
available.

Alternatively, it might occur due to scarcity. For example, the
exhaustion of IPv4 addresses creates a power differential between
those who have addresses and those who do not, which can affect how
the protocols that depend on IP connectivity are deployed and used.
If those addresses are held by only a few consolidation is the
result, in turn leading to ossification because new applications
cannot be deployed without their cooperation.

Some effects of inherited centralization can be mitigated by
enforcing layer boundaries using techniques like encryption. When
the number of parties who have access to content of communication
are limited, parties at lower layers can be prevented from
interfering with and observing it. Although those lower-layer

parties might still prevent communication, encryption also makes it more difficult to discriminate a target from other users' traffic.

Note that the prohibitive effect of encryption on inherited centralization is most pronounced when most (if not all) traffic is encrypted. See also [RFC7258].

2.2.5. Platform Centralization

The complement to inherited centralization is platform centralization -- where a function does not directly define a central role, but could facilitate consolidation in the applications it supports.

For example, HTTP [HTTP] is not generally considered a centralized protocol; interoperable servers are easy to instantiate, and multiple clients are available. It can be used without central coordination beyond that provided by DNS, as discussed above. However, applications built on top of HTTP (as well as the rest of the "Web Platform") often exhibit consolidation (for example, social networking). HTTP is therefore an example of platform centralization -- while the protocol itself is not centralized, it facilitates the creation of consolidated services and applications.

Like concentration, platform centralization is difficult to prevent with protocol design. Because of the layered nature of the Internet, most protocols allow considerable flexibility in how they are used, to promote permissionless innovation.

3. Decentralization

While the term "decentralization" has a long history of use in economics, politics, religion, and international development, Baran gave one of the first definitions relevant to computer networking, as a condition when "complete reliance upon a single point is not always required." [BARAN]

This seemingly straightforward technical definition hides several issues.

First, identifying which aspects of a function to decentralize and how to do so can be difficult, both because there are often many ways a function might be centralized, and because consolidation sometimes only becomes clear after the function is deployed at scale. Efforts to decentralize often have the effect of merely shifting consolidation to a different place.

For example, a cloud storage function might be implemented using a distributed consensus protocol, assuring that the failure of any single node will not affect the system's operation or availability.

In that sense, it is decentralized. However, if it is operated by a single legal entity, that brings a very different kind of centralization, especially if there are few other options available, or if there is friction against choosing other options.

Another example is the Web, which was envisioned and widely held to be a decentralizing force in its early life. Its inherent platform centralization only became apparent when large sites successfully leveraged network effects for dominance of social networking, marketplaces, and similar functions.

Second, different parties might have good-faith differences on what "sufficiently decentralized" means based upon their beliefs, perceptions and goals. Just as centralization is a continuum, so is decentralization, and not everyone agrees one what the "right" level or type is, how to weigh different forms of centralization against each other, or how to weigh potential consolidation against other architectural goals (such as security or privacy).

These tensions can be seen, for example, in the DNS. While much of the system is decentralized through the distribution of the lookup function to local servers that users have the option to override, the DNS is also a name space -- a single, global "source of truth" with inherent (if beneficial) centralization in its management. ICANN mitigates the associated risk through multi-stakeholder governance (see [Section 3.1.2](#)). While many believe that this arrangement is sufficient and might even have desirable qualities (such as the ability to impose community standards over the operation of the name space), others reject ICANN's oversight of the DNS as illegitimate, favoring decentralization based upon distributed consensus protocols rather than multistakeholderism. [[MUSIANI](#)]

Third, decentralization unavoidably involves adjustments to the power relationships between protocol participants, especially when it opens up the possibility of consolidation elsewhere. As Schneider notes in [[SCHNEIDERb](#)], decentralization "appears to operate as a rhetorical strategy that directs attention toward some aspects of a proposed social order and away from others", so "we cannot accept technology as a substitute for taking social, cultural, and political considerations seriously." Or, as more bluntly stated in [[BODO](#)], "without governance mechanisms in place, nodes may collude, people may lie to each other, markets can be rigged, and there can be significant cost to people entering and exiting markets."

For example, while blockchain-based cryptocurrencies might address the consolidation inherent in traditional currencies through technical means, the concentration of power that many exhibit in terms of voting/mining power, distribution of funds, and diversity

of codebase causes some to question how decentralized they actually are. [AREWEDECENTRALIZEDYET] The lack of formal structures brings an opportunity for latent, informal power structures that have their own risks -- including consolidation. [FREEMAN]

In practice, this means that decentralizing a function requires considerable work, is inherently political, and involves a large degree of uncertainty about the outcome. In particular, if one considers decentralization as a larger social goal (in the spirit of how the term is used in other, non-computing contexts), merely rearranging technical functions may lead to frustration. "A distributed network does not automatically yield an egalitarian, equitable or just social, economic, political landscape." [BODO]

3.1. Decentralization Techniques

Over time, a few different techniques have been used to facilitate decentralization of Internet protocols. The subsection below examine some of these techniques, along with their limitations.

None of them is a panacea; it is not possible to completely remove all forms of centralization from protocols that, at their heart, require agreement between multiple parties. However, when performed properly, decentralization might produce an outcome where that risk is understood, acceptable, and, where possible and appropriate, mitigated.

There is also no "correct" way to decentralize; it does not require that provision of a function need be distributed in a particular fashion, or to a particular degree. For example, the Domain Name System [RFC1035] is widely agreed to have an acceptable form of centralization, despite it being provided by a limited set of entities.

3.1.1. Federation

A common technique for addressing centralization in Internet protocols is federation -- designing them in such a way that new instances of a function are easy to create and can maintain interoperability and connectivity with other instances.

For example, SMTP [RFC5321] is the basis of the e-mail suite of protocols, which has two functions that exhibit centralization:

1. Giving each user a globally unique address, and

2. Routing messages to the user, even when they change network locations or become disconnected for long periods of time.

E-mail reuses DNS to help mitigate the first. To mitigate the
second, it defines a specific role for routing users' messages, the
Message Transfer Agent (MTA). By allowing anyone to deploy an MTA
and defining rules for interconnecting them, the protocol's users
avoid a requirement for a single central router.

Users can (and often do) choose to delegate that role to someone
else, or run their own MTA. However, many now consider running a
personal MTA to be impractical because of the likelihood of a small
MTA being classified as a spam source. Because large MTA operators
are widely known and have greater impact if their operation is
affected, they are less likely to be classified as such,
concentrating the protocol's operation (see Section 2.2.3).

Another example of a federated Internet protocol is XMPP [RFC6120],
supporting "instant messaging" and similar functionality. Like e-
mail, it reuses DNS for naming and requires federation to facilitate
rendezvous of users from different systems.

While some deployments of XMPP do support truly federated messaging
(i.e., a person using service A can interoperably chat with someone
using service B), many of the largest do not. Because federation is
voluntary, some operators captured their users into a single
service, denying them the benefits of global interoperability.

The examples above illustrate that, while federation can be a useful
technique for avoiding proprietary centralization and managing
beneficial centralization, it does not prevent concentration or
platform centralization. If a single entity can capture the value
provided by a protocol, it may use the protocol as a platform to get
a "winner take all" outcome -- a significant risk with many Internet
protocols, since network effects often promote such outcomes.
Likewise, external factors (such as spam control) might naturally
"tilt the table" towards a few operators.

3.1.2. Multi-Stakeholder Governance

Protocol designers sometime attempt to mitigate beneficial
centralization (see Section 2.2.2) by delegating that function's
governance to a multi-stakeholder body -- an institution that
includes representatives of the different kinds of parties that are
affected by the system's operation ("stakeholders") in an attempt to
make well-reasoned, legitimate, and authoritative decisions.

The most widely studied example of this technique is the governance
of the DNS name space, which as a "single source of truth" exhibits
beneficial centralization. The associated risk is managed through
administration by the Internet Corporation for Assigned Names and

Numbers (ICANN), a global multi-stakeholder body with representation
from end users, governments, operators, and others.

Another example is the governance of the Web's trust model,
implemented by Web browsers as relying parties and Certificate
Authorities as trust anchors. To promote the operational and
security requirements necessary to provide the desired properties,
the CA/Browser Forum was established as an oversight body that
involves both parties as stakeholders.

Yet another example of multi-stakeholderism is the standardization
of Internet protocols themselves. Because a specification controls
implementation behavior, the standardization process can be seen as
a single point of control. As a result, Internet standards bodies
like the IETF allow open participation and contribution, make
decisions in an open and accountable way, have a well-defined
process for making (and when necessary, appealing) decisions,
considering the views of different stakeholder groups [RFC8890].

A major downside of this approach is that setup and ongoing
operation of multi-stakeholder bodies is not trivial. Additionally,
their legitimacy cannot be assumed, and may be difficult to
establish and maintain (see, e.g., [PALLADINO]). This concern is
especially relevant if the function being coordinated is broad,
complex, and/or contentious.

3.1.3. Distributed Consensus

Increasingly, distributed consensus technologies (such as the
blockchain) are touted as a solution to consolidation issues. A
complete survey of this rapidly changing area is beyond the scope of
this document, but we can generalize about its properties.

These techniques attempt to avoid centralization by distributing
functions to members of a sometimes large pool of protocol
participants. They typically guarantee proper performance of a
function using cryptographic techniques (often, an append-only
transaction ledger). A particular task's assignment to a node for
handling usually cannot be predicted or controlled.

Sybil attacks (where a party or coordinated parties cheaply create
enough protocol participants to affect how consensus is judged) are
a major concern for these protocols. They encourage diversity in the
pool of participants using indirect techniques, such as proof-of-
work (where each participant has to show significant consumption of
resources) or proof-of-stake (where each participant has some other
incentive to execute correctly).

Use of these techniques can create barriers to proprietary and
inherited centralization. However, depending upon the application in

question, both concentration and platform centralization are still possible.

Furthermore, a protocol or an application can use distributed consensus for some functions, but still be centralized elsewhere -- either because those functions cannot be decentralized (most commonly, rendezvous and global naming; see Section 2.2.2) or because the designer has chosen not to because of the associated costs and lost opportunities.

Even when distributed consensus is used for all technical functions of a service, some coordination is still necessary -- whether that be through governance of the function itself, creation of shared implementations, or documentation of shared wire protocols. That represents centralization, just at a different layer (inherited or platform). For example, the Ethereum "merge" demonstrated that the blockchain could address environmental concerns, but only through coordinated community effort and governance. [ETHEREUM]

These potential shortcomings do not rule out the use of distributed consensus technologies in every instance. They do, however, caution against uncritically relying upon these technologies to avoid or mitigate centralization.

4. What Should Internet Standards Do?

Centralization is driven by powerful forces -- both economic and social -- as well as the network effects that come with Internet scale. Bodies like the IETF create voluntary standards; they cannot require adoption, but when a specification succeeds it creates those very network effects. As such, standards bodies cannot prevent all forms of consolidation, but they can still take meaningful steps to prevent some of them. The subsections below suggest a few.

4.1. Bolster Legitimacy

Many efforts to address Internet consolidation are likely to take place outside of standards bodies. If the IETF wishes to contribute to these efforts and assure their compatibility with the Internet's architectural goals, it must be seen as legitimate by the relevant parties -- especially, by competition regulators. Otherwise, if the IETF is perceived as representing or being controlled by "big tech" concerns, its ability to guide decisions that affect the Internet will be diminished considerably.

The IETF already has features that arguably provide considerable legitimacy; for example, open participation and representation by individuals rather than companies both enhance input legitimacy; a well-defined process with multiple layers of appeals and transparency contributes to throughput legitimacy, and a long

history of successful Internet standards provides perhaps the
strongest source of legitimacy for the IETF -- its output.

However, it is also widely recognized the considerable costs (not
just financial) involved in successfully participating in the IETF
have a tendency to favour larger companies over smaller concerns.
Likewise, the specialised and highly technical nature of the work
creates barriers to entry for non-technical stakeholders. These
factors have the potential to reduce the legitimacy of the IETF's
decisions, at least in some eyes.

Efforts to address these shortcomings are ongoing; see, for example,
[RFC8890]. Overall, bolstering the legitimacy of the organization
should be seen as a continuous effort.

4.2. Engage with Centralization Thoroughly but Realistically

Some kinds of centralization are easy to manage in standards
efforts. For example, if a protocol with a fixed role for a single
party were to be proposed to the IETF for publication as a standard,
it would be rejected out of hand. There is a growing body of
knowledge and experience in managing the risks of beneficial
centralization, and a strong inclination to reuse existing
infrastructure where possible. As discussed above, encryption is
often a way to manage inherited centralization, and has become the
norm in standard protocols. These responses are appropriate ways for
Internet standards to manage centralization.

However, mitigating concentration and platform centralization is
much more difficult in standards efforts. Because the IETF has no
"protocol police", it's not possible to demand, for example, that
someone stop building a proprietary service using a federated
protocol; even if it could, doing so would contradict architectural
goals like permissionless innovation. While the imprimatur of an
Internet Standard is not without value, merely withholding it cannot
prevent these sources of consolidation.

Therefore, committing significant resources to scrutinizing
protocols for latent centralization -- especially for concentration
and platform centralization -- is unlikely to be effective in
preventing Internet consolidation. Almost all existing Internet
protocols -- including IP, TCP, HTTP, and DNS -- exhibit
concentration or platform centralization. Refusing to standardize a
newer protocol because it exhibits similar properties would not be
equitable, proportionate, or effective.

When claims are made that a given proposal is "centralized" or
"decentralized", the context of those statements should be examined
for presuppositions, assumptions, and omissions. One framework for

critical interrogations is offered by [BACCHI], which can be adapted
for centralization-related discussions:

1. What is the nature of the centralization that is represented as
   being problematic?

2. What deep-seated presuppositions or assumptions (conceptual
   logics) underlie this representation of the "problem"?

3. How has this representation of the problem come about?

4. What is left unproblematic in this problem representation?
   Where are the silences? Can the "problem" be conceptualized
   differently?

5. What effects are produced by this representation of the
   "problem"?

6. How and where has this representation of the "problem" been
   produced, disseminated and defended? How has it been and/or how
   can it be disrupted and replaced?

[SCHNEIDERa] implores that proposals to decentralize be "really,
really clear about what particular features of a system a given
design seeks to decentralize" and promotes borrowing remedies from
more traditional governance systems, such as separation of powers
and accountability.

When centralization is found, standards efforts should consider its
relationship with architectural goals as they consider how to
address it. In particular, attention should be paid to how effective
standards (as a form of architectural control) is in achieving each
goal.

For example, privacy is often more effectively ensured by ex ante
technical constraints, as compared to ex post legal regulation.
Conversely (as discussed) some centralization may be more
effectively addressed through legal regulation. Thus, a standards
effort balancing these concerns might focus primarily on privacy.
However, these are often not completely separable goals.
Concentration can result in one or a few entities having greater
volume and variety of data available exclusively to them, raising
significant privacy and security concerns.

4.3. Decentralize Proprietary Functions

It is worthwhile to create specifications for functions that are
currently only available from proprietary providers. Open standards
can provide a viable alternative to a consolidated function.

Such efforts might include large-scale protocols for existing proprietary functions (e.g., chat) as well as smaller efforts to improve interoperability and portability of specific features that are often used to "lock in" users to a platform; for example, a format for lists of contacts in a social network.

A common objection to this approach is that adoption is voluntary, not mandatory; there are no "standards police" to mandate their use or enforce correct implementation. For example, specifications like [ACTIVITYSTREAMS]) have been available for some time without being used in a federated manner by commercial social networking providers.

However, while standards aren't mandatory, legal regulation is, and legal mandates for interoperability are increasingly discussed by policymakers as a remedy for competition issues (see, e.g., [OECD]).

As such, appetite for regulation presents an opportunity for new specifications to decentralize these functions, backed by a legal mandate in combination with changing norms and the resulting market forces [LESSIG]. That opportunity also presents a risk, however, if the resulting legal regulation is at odds with the Internet architecture.

Successfully creating standards that work in concert with legal regulation presents many potential pitfalls, and will require improved and new capabilities (especially liaison, likely originating in the IAB), and considerable effort. If the Internet community does not make that effort, it is likely that regulators will turn to other sources of interoperability specifications -- most likely, with less transparency, more narrow input, limited experience, and without reference to the Internet's architectural goals.

4.4. Evaluate New Decentralization Techniques

The decentralization techniques listed in Section 3.1 are not a closed set; wide interest has spurred development of new approaches, both in general and as solutions to specific problems.

For example, secure multi-party computation techniques (see, e.g., [YAO]) can be composed to allow parties to compute over private inputs without revealing them. Protocols like [ENPA] and [PRIO] use them to limit the information available to participants in protocols to realize privacy goals; as discussed in Section 4.6 doing so might also counteract some sources of centralization. However, as with those covered above, these techniques do not automatically preclude all consolidation; such systems often still require trust, even if

it is limited, and that might result in other forms of consolidation emerging.

Whether use of these techniques (or others) can meaningfully counteract consolidation is still uncertain. Standards bodies (including the IETF) can serve an important function by incubating them, applying (and, where necessary, developing) architectural guidelines for privacy, security, operability, and other goals, and assuring interoperability. When appropriate, publication on the standards track or as experimental can send signals to implementers, users, and regulators about their fitness for particular purposes.

4.5. Enable Switching

To minimize centralization, standards-defined functions should have an explicit goal enabling users' switching between implementations and deployments of protocols.

One necessary condition for this is the availability of alternatives; breadth and diversity of implementation and deployment are required. Section 2.1 of [RFC5218] explores some factors in protocol design that encourage this outcome.

Another factor is the cost of substituting an alternative implementation or deployment by users. This implies that the standard needs to be functionally complete and specified precisely enough to allow substitution.

These goals of completeness and diversity are sometimes in tension. If a standard becomes extremely complex, it may discourage implementation diversity because the cost of a complete implementation is too high (consider: Web browsers). On the other hand, if the specification is too simple, it may not enable easy switching, especially if proprietary extensions are necessary to complete it (see Section 4.7).

One objection to protocols that enable easy switching is that they reduce the incentives for implementation by commercial vendors. While a completely commoditized protocol might not allow implementations to differentiate themselves, they provide opportunities for specialization and improvement elsewhere in the value chain [CHRISTENSEN]. Well-timed standards efforts leverage these forces to focus proprietary interests on top of open technology, rather than as a replacement for it.

4.6. Control Delegation of Power

Some functions might see substantial benefits if they are provided
by a third party in communication. Adding a new party to
communication can improve:

   Efficiency: Many functions on the Internet are more efficient
    when performed at a higher scale. For example, a content delivery
    network can offer services at a fraction of the financial and
    environmental cost that someone serving content themselves would
    otherwise pay, because of the scale they operate at. Likewise, a
    two-sided market platform can introduce sizeable efficiencies
    over pair-wise buyer/seller trading [SPULBER].

   Simplicity: Completely disintermediating communication can shift
    the burden of functions onto endpoints. This can cause increased
    cognitive load for users; for example, compare commercial social
    networking platforms with self-hosted efforts.

   Specialization: Having a function concentrated into a few hands
    can improve outcomes because of the resulting specialization. For
    example, services overseen by professional administrators are
    often seen to have a better security posture and improved
    availability.

   Privacy: For some functions, user privacy can be improved by
    concentrating their activity to prevent individual behaviors from
    being discriminated from each other.[CHAUM] Introduction of a
    third party can also enforce functional boundaries -- for
    example, to reduce the need for users to trust potentially
    malicious endpoints, as seen in the so-called "oblivious"
    protocols (e.g., [RFC9230]) that allow end users to hide their
    identity from services, while still accessing them.

However, introducing a new party to communication adds concentration
and platform centralization to Internet functions, because it brings
opportunities for control and observation. While (as discussed
above) standards efforts have a very limited capability to prevent
or control the resulting consolidation, designing functions with
thoughtful constraints on third party functions can prevent at least
the most egregious outcomes.

Most often, third parties are added to functions as "intermediaries"
or in designated "agent" roles. In general, they should only be
interposed because of the positive action of at least one of the
primary parties, and should have their ability to observe or control
communication limited to what is necessary to perform their intended
function.

For example, early deployments of HTTP allowed intermediaries to be interposed by the network without knowledge of the endpoints, and those intermediaries could see and change the full content of traffic by default -- even when they are only intended to perform basic functions such as caching. Because of the introduction of HTTPS and the CONNECT method (see Section 9.3.6 of [HTTP]), combined with efforts to encourage its adoption, those intermediaries are now required to be explicitly interposed by one endpoint.

See [I-D.thomson-tmi] for more guidance on protocol intermediation.

The term "intermediary" is also used (often in legal and regulatory contexts) more broadly than it has been in protocol design; for example, an auction Web site intermediates between buyers and sellers is considered an intermediary, even though it is not formally an intermediary in HTTP (see Section 3.7 of [HTTP]). Protocol designers can address the centralization associated with this kind of intermediation by standardising the function, rather than restricting the capabilities of the underlying protocols; see Section 4.3.

4.7. Consider Extensibility and Modularity Carefully

The Internet's ability to evolve is critical, allowing it to meet new requirements and adapt to new conditions without requiring a "flag day" to upgrade implementations. Typically, functions accommodate evolution by defining extension interfaces, which allow optional features to be added or change over time in an interoperable fashion.

However, these interfaces can also be a basis for platform centralization if a powerful entity can change the target for meaningful interoperability by adding proprietary extensions to a standard. This is especially true when the core standard does not itself provide sufficient utility on its own.

For example, the SOAP protocol's [SOAP] extreme flexibility and failure to provide significant standalone value allowed vendors to require use of their preferred extensions, favoring those who had more market power.

Therefore, standards efforts should focus on providing concrete utility to the majority of their users as published, rather than being a "framework" where interoperability is not immediately available. Internet functions should not make every aspect of their operation extensible; boundaries between modules should be designed in a way that allows evolution and discourages consolidation, while still offering meaningful functionality.

Beyond allowing evolution, well-considered interfaces can also aid decentralization efforts. The structural boundaries that emerge between the sub-modules of the function -- as well as those with adjacent functions -- provide touchpoints for interoperability and an opportunity for substitution of providers.

In particular, if the interfaces of a function are well-defined and stable, there is an opportunity to use different providers for that function. When those interfaces are open standards, change control resides with the Internet community instead of remaining in proprietary hands, further enhancing stability and enabling (but not ensuring) decentralization.

5. Security Considerations

This document does not have a direct security impact on Internet protocols. However, failure to consider centralization might cause a myriad of security issues.

6. Informative References

[ABRAHAMSON] Abrahamson, Z., "Essential Data", Yale Law Journal, Vol. 124, No. 3, 2014, <https://www.yalelawjournal.org/comment/essential-data>.

[ACTIVITYSTREAMS] Prodromou, E., Ed. and J. Snell, Ed., "Activity Streams 2.0", W3C CR CR-activitystreams-core-20161215, W3C CR-activitystreams-core-20161215, 15 December 2016, <https://www.w3.org/TR/2016/CR-activitystreams-core-20161215/>.

[ALIGIA] Aligia, P. D. and V. Tarko, "Polycentricity: From Polanyi to Ostrom, and Beyond", Governance: An International Journal of Policy, Administration, and Institutions, Vol. 25, No. 2, p. 237, April 2012, <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-0491.2011.01550.x>.

[AREWEDECENTRALIZEDYET] bitcoinera, "Are We Decentralized Yet?", 2022, <https://bitcoinera.app/arewedecentralizedyet/>.

[BACCHI] Bacchi, C., "Introducing the 'What's the Problem Represented to be?' approach", Chapter 2, Engaging with Carol Bacchi, 2012, <https://library.oapen.org/bitstream/handle/20.500.12657/33181/560097.pdf?sequence=1#page=34>.

[BARABASI] Albert, R., "Emergence of Scaling in Random Networks", SCIENCE, Vol. 286, No. 15, p. 509, October 1999, <https://barabasi.com/f/67.pdf>.

[BARAN]   Baran, P., "On Distributed Communications: Introduction
             to Distributed Communications Networks", 1964, <https://
             www.rand.org/pubs/research_memoranda/RM3420.html>.

[BCP95]   Alvestrand, H., "A Mission Statement for the IETF", BCP
             95, RFC 3935, October 2004.

[BODO]   Bodó, B., Brekke, J. K., and J.-H. Hoepman,
             "Decentralization: a multidisciplinary perspective",
             Internet Policy Review, Vol. 10, No. 2, June 2021,
             <https://doi.org/10.14763/2021.2.1563>.

[CHAUM]   Chaum, D. L., "Untraceable Electronic Mail, Return
             Addresses, and Digital Pseudonyms", Communications of the
             ACM, Vol. 24, No. 2, February 1981, <https://dl.acm.org/
             doi/10.1145/358549.358563>.

[CHRISTENSEN] Christensen, C., "The Law of Conservation of
             Attractive Profits", Harvard Business Review,
             "Breakthrough Ideas for 2004", February 2004.

[ECH]   Rescorla, E., Oku, K., Sullivan, N., and C. A. Wood, "TLS
             Encrypted Client Hello", Work in Progress, Internet-
             Draft, draft-ietf-tls-esni-15, 3 October 2022, <https://
             datatracker.ietf.org/doc/html/draft-ietf-tls-esni-15>.

[ENPA]   Apple and Google, "Exposure Notification Privacy-
             preserving Analytics (ENPA) White Paper", April 2021,
             <https://covid19-static.cdn-apple.com/applications/
             covid19/current/static/contact-tracing/pdf/
             ENPA_White_Paper.pdf>.

[ETHEREUM]   Ethereum, "The Merge", February 2023, <https://
             ethereum.org/en/upgrades/merge/>.

[FARRELL]   Farrell, H. and A. L. Newman, "Weaponized
             Interdependence: How Global Economic Networks Shape State
             Coercion", International Security, Vol. 44, No. 1, p. 42,
             2019, <https://doi.org/10.1162/ISEC_a_00351>.

[FREEMAN]   Freeman, J., "The Tyranny of Structurelessness", Berkeley
             Journal of Sociology, Vol. 17, 1972, <https://
             www.jstor.org/stable/41035187>.

[HTTP]   Fielding, R., Ed., Nottingham, M., Ed., and J. Reschke,
             Ed., "HTTP Semantics", STD 97, RFC 9110, DOI 10.17487/
             RFC9110, June 2022, <https://www.rfc-editor.org/rfc/
             rfc9110>.

[I-D.thomson-tmi]
              Thomson, M., "Principles for the Involvement of
              Intermediaries in Internet Protocols", Work in Progress,
              Internet-Draft, draft-thomson-tmi-04, 8 September 2022,
              <https://datatracker.ietf.org/doc/html/draft-thomson-
              tmi-04>.

[ISOC]   Internet Society, "Consolidation in the Internet
              Economy", Internet Society Global Internet Report, 2019,
              <https://future.internetsociety.org/2019/>.

[JUDGE]   Judge, K., "Intermediary Influence", University of
              Chicago Law Review, Vol. 82, p. 573, 2014, <https://
              scholarship.law.columbia.edu/faculty_scholarship/1856>.

[KENDE]   Kende, M., Kvalbein, A., Allford, J., and D. Abecassis,
              "Study on the Internet's Technical Success Factors",
              December 2021, <https://blog.apnic.net/wp-content/
              uploads/2021/12/MKGRA669-Report-for-APNIC-LACNIC-V3.pdf>.

[LESSIG]   Lessig, L., "The New Chicago School", Journal of Legal
              Studies, Vol. 27, June 1998, <https://
              www.journals.uchicago.edu/doi/10.1086/468039>.

[MADISON]   Madison, J., "The Structure of the Government Must
              Furnish the Proper Checks and Balances Between the
              Different Departments", The Federalist Papers, No. 51,
              February 1788.

[MOXIE]   Marlinspike, M., "Reflections: The ecosystem is moving",
              May 2016, <https://signal.org/blog/the-ecosystem-is-
              moving/>.

[MUSIANI]   Musiani, F., "Alternative Technologies as Alternative
              Institutions: The Case of the Domain Name System", The
              Turn to Infrastructure in Internet Governance, 2016,
              <https://link.springer.com/chapter/
              10.1057/9781137483591_4>.

[OECD]   OECD, "Data portability, interoperability and digital
              platform competition", June 2021, <https://www.oecd.org/
              daf/competition/data-portability-interoperability-and-
              digital-platform-competition-2021.pdf>.

[OLUMOFIN]   Olumofin, F. and I. Goldberg, "Revisiting the
              Computational Practicality of Private Information

Retrieval", 2010, <https://link.springer.com/chapter/
10.1007/978-3-642-27576-0_13>.

[PALLADINO] Palladino, N. and N. Santaniello, "Legitimacy, Power,
             and Inequalities in the Multistakeholder Internet
             Governance", 2020, <https://link.springer.com/book/
             10.1007/978-3-030-56131-4>.

[PRIO]   Corrigan-Gibbs, H. and D. Boneh, "Prio: Private, Robust,
             and Scalable Computation of Aggregate Statistics", March
             2017, <https://crypto.stanford.edu/prio/paper.pdf>.

[RFC1035]   Mockapetris, P., "Domain names - implementation and
             specification", STD 13, RFC 1035, DOI 10.17487/RFC1035,
             November 1987, <https://www.rfc-editor.org/rfc/rfc1035>.

[RFC4271]   Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A
             Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI

10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/rfc/rfc4271>.

[RFC5218]   Thaler, D. and B. Aboba, "What Makes for a Successful Protocol?", RFC 5218, DOI 10.17487/RFC5218, July 2008, <https://www.rfc-editor.org/rfc/rfc5218>.

[RFC5321]   Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <https://www.rfc-editor.org/rfc/rfc5321>.

[RFC6120]   Saint-Andre, P., "Extensible Messaging and Presence Protocol (XMPP): Core", RFC 6120, DOI 10.17487/RFC6120, March 2011, <https://www.rfc-editor.org/rfc/rfc6120>.

[RFC7258]   Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <https://www.rfc-editor.org/rfc/rfc7258>.

[RFC791]   Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/rfc/rfc791>.

[RFC793]   Postel, J., "Transmission Control Protocol", RFC 793, DOI 10.17487/RFC0793, September 1981, <https://www.rfc-editor.org/rfc/rfc793>.

[RFC8890]   Nottingham, M., "The Internet is for End Users", RFC 8890, DOI 10.17487/RFC8890, August 2020, <https://www.rfc-editor.org/rfc/rfc8890>.

[RFC9230]   Kinnear, E., McManus, P., Pauly, T., Verma, T., and C.A. Wood, "Oblivious DNS over HTTPS", RFC 9230, DOI 10.17487/RFC9230, June 2022, <https://www.rfc-editor.org/rfc/rfc9230>.

[SCHNEIDERa] Schneider, N., "What to do once you admit that decentralizing everything never seems to work", Hacker Noon, October 2022, <https://nathanschneider.info/articles/DecentralHacker.html>.

[SCHNEIDERb] Schneider, N., "Decentralization: an incomplete ambition", Journal of Cultural Economy, Vol. 12, No. 4, 2019, <https://osf.io/m7wyj/>.

[SOAP]   Mitra, N., Ed. and Y. Lafon, Ed., "SOAP Version 1.2 Part 0: Primer (Second Edition)", W3C REC REC-soap12-part0-20070427, W3C REC-soap12-part0-20070427, 27 April 2007, <https://www.w3.org/TR/2007/REC-soap12-part0-20070427/>.

[SPULBER]
Spulber, D. F., "Solving The Circular Conundrum: Communication And Coordination In Internet Markets", Northwestern University Law Review, Vol. 104, No. 2, 2010, <https://wwws.law.northwestern.edu/research-faculty/clbe/workingpapers/documents/spulber_circularconundrum.pdf>.

[YAO]   Yao, A. C., "Protocols for secure computations", SFCS '82, November 1982, <https://dl.acm.org/doi/10.5555/1382436.1382751>.

Appendix A. Acknowledgements

This document benefits from discussions with Brian Trammell during our shared time on the Internet Architecture Board.

Thanks to Jari Arkko, Kristin Berdan, Christian Huitema, Mallory Knodel, Eliot Lear, Tommy Pauly, and Martin Thomson for their comments and suggestions.

No machine learning models were used in the production of this document.

Author's Address

Mark Nottingham
Prahran
Australia

Email: mnot@mnot.net
URI: https://www.mnot.net/