Workgroup: Network Working Group
Internet-Draft:
draft-nottingham-avoiding-internet-
centralization-06
Published: 3 December 2022
Intended Status: Informational
Expires: 6 June 2023
Authors: M. Nottingham

Internet Consolidation: What can Standards Efforts Do?

Abstract

Despite the Internet being designed and operated as a decentralized network-of-networks, forces continuously emerge to encourage and sometimes enforce consolidation of power into few hands.

This document offers a definition of consolidation and relates it to centralization, explains why they 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. 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 positioned as a global public good because joining, deploying an application on, or using the Internet does not require permission from or ceding control to a single entity.

While avoiding consolidation of power on the Internet remains a widely shared goal, achieving it consistently has proven difficult. Today, many successful protocols and applications on the Internet operate in a centralized fashion -- to the point where some proprietary services have become so well-known that they are commonly mistaken for the Internet itself. Even when 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 performed 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 that relate to Internet standards efforts, and argues that while the IETF may not be able to prevent consolidation, there are still meaningful steps we can take to counteract it.

Section 2 defines consolidation and centralization, explains why and when they are undesirable, and surveys contributors to consolidation seen 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 consolidation and mitigating its effects.

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

2. Consolidation and Centralization

This document defines "consolidation" as the ability of a single entity or a small group of them to exclusively observe, capture, control, or extract rent from the operation or use of an Internet function.

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 defined broadly. 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.

However, the Internet's 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 the forces of consolidation, as can supply of transit between networks (so called "Tier 1" networks).

"Centralization" measures the contribution of a function's technical design to consolidation. As such, it is a primarily architectural phenomenon. For example, many consider the social networking market to be highly consolidated around a few providers; the technologies that they use are proprietarily centralized (see Section 2.2.1) and thus contribute to that consolidation.

Centralization is not a binary condition; a function's design might contribute to or be vulnerable to consolidation in multiple ways and various degrees. Even when decentralization techniques are purposefully used to avoid it, centralization often appears in other aspects of the function's design -- for example, in its governance, implementation, deployment, or in ancillary functions. As Schneider says, "decentralized technology alone does not guarantee decentralized outcomes." [SCHNEIDERb]

Therefore, this document considers the amount of "consolidation risk" associated with a function's design, depending on the scale, scope, and nature of those contributions and vulnerabilities.

2.1. Assessing Consolidation Risk

By default, Internet protocol designers avoid centralized designs, because the Internet's very nature is incompatible with centralization. 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. [SCHNEIDERa] 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."

With that in mind, consolidation risk 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, consolidation risk 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 these functions' implementation leads to a more robust outcome when viewed systemically. [ALIGIA]

* *Self-Reinforcement*: As widely noted (see, e.g., [VESTAGER]), 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 need to be evaluated carefully on a case-by-case basis.

For example, it is important to distinguish consolidation risk 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 centralization, only courts (and in some cases, 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, and 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 due to 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. Furthermore, a failure due to 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, consolidation 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. Contributors to Centralization

A function's design can exhibit centralization in a variety of ways. The subsections below describe different contributors to and expressions of centralization in Internet functions.

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, they have corresponding consolidation risk -- 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 the introduction of a centralized function. In doing so, they are explicitly relying on centralization to deliver a particular benefit.

For example, a function that needs 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.

Another function exhibiting beneficial centralization is IP addresses allocation. Internet routing requires addresses to be allocated uniquely, but if a single government or company captured the addressing function, the entire Internet would be at risk of abuse by that entity. Similarly, the need for coordination in the Web's trust model brings consolidation risk, because of 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 has consolidation risk.

A centralized function's inherent power can also be used to beneficial ends. For example, when traffic from many users is mixed together 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 consolidation 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 consolidation risk. One commonly seen application of this kind of beneficial centralization is in content moderation functions.

When beneficial centralization is present, Internet protocols often attempt to mitigate the associated risks using measures such as federation (see Section 3.1.1) and multi-stakeholder governance (see Section 3.1.2). Protocols that successfully mitigate the associated consolidation risks 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 proprietary centralization and mitigates any beneficial centralization present, 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." Economic, legal, and social factors that encourage use of a central function despite the absence of such a requirement in the protocol itself can cause concentration.

Often, the factors driving concentration are related to 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, network 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, consolidation risk is indicated. Conversely, a function based on a well-defined, open specification designed to minimize switching costs might be considered to have less consolidation risk even when there are only a few large providers.

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, a group of people who wish to communicate are often locked in by the choices that their peers make, 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" protocols 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 consolidation risk 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 inability) 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 an inherited consolidation risk, because applications that use it are vulnerable to the control it has over their operation. Even Open Source projects can exhibit this risk if there are factors that make forking difficult (for example, the cost of maintaining that fork).

Inherited centralization 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.

Some kinds of inherited centralization can be prevented by enforcing layer boundaries through use of 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 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 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, often in a way that it becomes attractive to form a dependency on one party's operation.

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 evident after the function is deployed at scale.

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 consolidation risk, 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 consolidation is a continuum, so is

decentralization, and not everyone agrees one what the "right" level or type is, how to weigh different forms of consolidation against each other, or how to weigh 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 of its management. The associated risk is mitigated through multi-stakeholder governance by ICANN (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 decentralizing a function opens up the possibility of consolidation elsewhere. As Schneider notes in [SCHNEIDERa], 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

In the context of Internet standardization, decentralization is implemented as a two-step process: assessing the nature of consolidation risk, followed by the application of techniques to reduce or mitigate it. The subsections below examine some of these techniques.

Choosing the appropriate techniques for decentralization requires balancing the specific goals of the function against consolidation risk, because completely precluding all forms of consolidation through technical means is rarely achievable. When performed properly, decentralization might produce an outcome that still has consolidation risk, but that risk should be understood, acceptable, and, where possible and appropriate, mitigated.

Notably, decentralization 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 acceptable consolidation risk, despite it being provided by a limited set of entities.

3.1.1. Federation

A widely known technique for managing consolidation 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 have consolidation risk:

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 risk. 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, running your own mail server has become difficult, 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 to avoid proprietary centralization and manage 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 mitigate the consolidation risks associated with a beneficial centralized function (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 mitigated 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 ensure that all parties meet 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 consolidation risk 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.

It is also important to recognize that a protocol or an application can use distributed consensus for some functions, but still have consolidation risk 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 consolidation risk, just at a different layer (inherited or platform).

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 consolidation.

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. Because permissionless innovation is a core value for the Internet, and yet much of the consolidation seen on the Internet is performed by proprietary platforms that take advantage of this nature, the controls available to standards efforts are very limited.

While standards bodies on their own cannot prevent consolidation, the subsections below suggest meaningful steps that can be taken.

4.1. Engage with Centralization Risk Thoroughly but Realistically

Some consolidation risks are easy to manage in standards efforts. For example, if a proprietary protocol were to be proposed to the IETF, it would be rejected out of hand. There is a growing body of knowledge and experience in managing the risk 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 consolidation risk.

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 that someone stop building a proprietary service using a federated protocol (for example). The standards process also cannot stop someone from building services "on top" of standard protocols without abandoning 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 consolidation risk -- especially for concentration and platform risks -- 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 faces similar risks 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. [BACCHI] offers one framework for such critical interrogations. [SCHNEIDERb] 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 consolidation risk is found, standards efforts should consider its relationship with other 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 consolidation risks may be more effectively addressed through legal regulation. Thus, as a first order concern, a standards effort balancing these concerns might focus primarily on privacy. However, often these are 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.

4.2. Decentralize Proprietary Functions

It is worthwhile to create specifications for functions that are currently only satisfied by proprietary providers. By building open specifications on top of already established standards, an alternative to a consolidated function can be created.

A common objection to such efforts 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 broad adoption by social networking providers.

However, while standards aren't mandatory, legal regulation is, and regulators around the globe are focusing specific efforts on some aspects of the Internet. In particular, legal mandates for interoperability are increasingly discussed as a remedy for competition issues (see, e.g., [OECD]).

As such, appetite for Internet regulation presents not just a risk to the Internet; it also constitutes 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 is new ground for the IETF, presents many potential pitfalls, and will require 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.

4.3. 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 4.5 doing so might also counteract some sources of centralization. However, as in other cases 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 signal implementers, users, and regulators about their fitness.

4.4. Build Robust Ecosystems

To minimize inherited consolidation risk, standards-defined functions should have an explicit goal of broad, diverse implementation and deployment so that users have as many choices as possible.

Section 2.1 of [RFC5218] explores some factors in protocol design that encourage this outcome.

This goal can also be furthered by ensuring that the cost of switching to a different implementation or deployment is as low as possible to facilitate subsequent substitution. This implies that the standard is functionally complete and specified precisely enough to result in meaningful interoperability.

The goals of completeness and diversity are sometimes in tension. If a standard is 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 offer enough functionality to be complete, and the resulting proprietary extensions may make switching difficult (see Section 4.6).

Also worthy of attention are the underlying incentives for implementation. 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.

Balancing these factors to build robust ecosystems is difficult, but is often helped by community building and good design -- in particular, appropriate use of layering. It also requires continuing maintenance and evolution of protocols, to assure that they are still relevant and appropriate to their use.

4.5. Control Delegation of Power

Some functions might see substantial benefits if they are performed by a third party in communication. When used well, 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 risk 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 consolidation risk associated with this kind of intermediation by standardising the function, rather than restricting the capabilities of the underlying protocols; see Section 4.2.

4.6. Consider Extensibility and Modularity Carefully

An important feature of the Internet is its ability to evolve, so that it can 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 consolidation risks might cause a myriad of security issues.

6. Informative References

[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://papers.ssrn.com/sol3/papers.cfm?abstract_id=2149165>.

[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. <https://www.rfc-editor.org/info/bcp95>

[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>.

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

[SCHNEIDERb] 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>.

[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>.

[VESTAGER] Vestager, M., "Defending Competition in a Digitised World, Address at the European Consumer and Competition Day", April 2019, <https://wayback.archive-it.org/12090/20191129202059/https://ec.europa.eu/commission/commissioners/2014-2019/vestager/announcements/defending-competition-digitised-world_en>.

[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.

Author's Address

Mark Nottingham
Prahran
Australia

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