Journal of Cyber Policy

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rcyb20

Guest Editorial

To cite this article: (2020) Guest Editorial, Journal of Cyber Policy, 5:1, 5-8, DOI: 10.1080/23738871.2020.1748081

To link to this article: https://doi.org/10.1080/23738871.2020.1748081

JOURNAL OF CYBER POLICY
2020, VOL. 5, NO. 1, 5–8
https://doi.org/10.1080/23738871.2020.1748081

Routledge
Taylor & Francis Group

Guest Editorial

By Andrew Sullivan, President and Chief Executive Officer, Internet Society

So what?
Why anyone should care about consolidation and the Internet

This issue of the *Journal of Cyber Policy* focusses on Internet consolidation, which can be described – in line with the definition that Jari Arkko uses in his article – as the 'increasing control over Internet infrastructure and services by a small set of organizations'. The Internet Society and Chatham House worked together to produce this issue, because we think the topic is a current, pressing one – and one that policymakers are beginning to act upon without always completely grasping the nuances.

There can be no question that consolidation is an important topic for all of us. As societies have become more dependent upon the Internet for communication and (for that matter) the elementary functioning of our social infrastructure, we need to care about how the Internet evolves. It does appear that the Internet is gradually being pulled together under the ownership and control of a small number of companies. In 2019, the Internet Society released the Global Internet Report, *Consolidation in the Internet Economy*, as a first attempt to map these trends, and to distil the implications for the Internet's future evolution. The topic is complicated, and the consequences unclear. That is why we welcomed the opportunity to collaborate with Chatham House on this special issue to encourage more in-depth research and thinking.

The reason the broad topic is complicated goes back to the nature of the Internet. Often, public policy initially approaches an issue on the Internet as just another telecommunications issue. This means using an approach designed for systems that the Internet has largely replaced. Now, there are two critical differences with the Internet that allowed it to become that replacement for other kinds of telecommunications. Each difference presents a barrier to taking old regulatory approaches and simply applying them to the Internet

First, because it is a network of networks, there is no centre and therefore no centre of control. This means that regulatory stances that start from the assumption of central control or policymaking are either ineffective or warp the development of networks away from the strengths of the Internet. It would be possible, for instance, to develop networks that connect together but that rely on greater central control. Such a development would not provide the same advantages to innovation as the global network of networks we have.

Second, the Internet mostly puts 'intelligence at the edge' – the network simply provides the transport of data, while the mechanisms to deal with the data are actually

handled by the network application, usually on individual devices. The network is agnostic about what kinds of applications it can host. So innovations mostly happen at places that are not really natural points for regulatory control. Often, indeed, the innovations happen before anyone (including any regulator) is aware that someone is inventing something.

The Internet is also often described in an idealised way that makes us imagine it is much more uniform than it is. I just committed that sin, above: 'a network of networks' makes it sound like each network is an atom that is roughly the same, and together they make a molecule that we call the Internet. In fact, however, the manifold networks differ enormously. Think of a defence department's network, a university's network, a small company's network, a multinational corporation's network, the network of a social media company and the network of a home Internet service provider. Together (and with other networks), they make up the Internet when they interoperate using open Internet protocols. But making regulations for 'the Internet' is quite likely to affect them in different ways, and something socially beneficial aimed at one of them could easily be quite harmful to another.

Yet this entire description also shows why extreme consolidation is bad. For part of what makes the Internet so strong is also why it can be seriously weakened by one party holding too much operational or political power. Consolidation may undermine the entire effort.

Indeed, this manifold nature shows that consolidation (or anyway, certain kinds of network concentration) is not always bad in every way, even while illustrating that it could be harmful. For instance, as Jesse Sowell describes in this issue, Internet exchange points ('IXPs') are a feature of the Internet, invented many years ago, to provide efficiency and lower costs. In extremely simple terms, IXPs make it possible for many different networks to connect to each other in a single place, instead of having to connect to each other one by one. This simplifies routing, makes interconnection faster, makes many Internet operations faster, and (in well-run IXPs) lowers the costs of networking for all the participating networks. None of that means that a single IXP per country or region (or for that matter a single global one) would be a good idea. That would make the whole system more brittle and subject to subversion.

Even in the case of particular services, consolidation can be beneficial. Consider that some kinds of Internet infrastructure used to be operated almost always directly by every network operator. Companies usually operated their own Domain Name System (DNS) servers both for serving answers and also for performing resolution. The same was true of email and even web servers. Everyone had system administrators who looked after the services. That mode of operation did not last, and everyone benefitted from the change.

The example services are all critical business systems, so they need to work no matter what. Yet ensuring perfect uptime can be expensive. Virtually anyone can operate a web server, for instance, with 90 per cent reliability ('one nine'). But that is more than 35 days a year of outage, and no business could afford to have their website down nearly three days every month. The industry standard is 'five nines' – 99.999% uptime. That represents just over five minutes of service unavailability in a year. It is expensive to have people available all day, every day in order to deal with trouble and ensure the systems operate reliably.

So, network operators naturally took to purchasing services from specialists. The specialists had many customers, none of whom wanted for themselves the kind of

dedicated staff that the specialist company could afford to hire. The network services became more reliable for everyone, everyone's costs went down, and reliability actually went up because even very small operators could afford to run a world-class operation. While this is just standard economic specialisation in action, it involves some market consolidation. Standard Internet functions are fundamentally commodities. Commodity markets tend to consolidate because of efficiency gains. At the same time, the dynamic has led to gains nobody really anticipated. Today, individuals and businesses can rent, by the hour, enormous computational power and huge quantities of storage – capabilities that approximately nobody in the world could have afforded to operate only a decade ago.

All of this comes with another issue. Many of the services that people rely upon on the Internet are deployed as web services. They depend on application program interfaces (APIs) that are exclusively controlled by the provider of those services. These are sometimes called 'walled gardens', but also form part of other applications' supply chain as discussed by both Riley and Cobbe et al. in this volume. Now, we might think that these are 'internet services' that are merely dependent on the Internet, not really part of it. Yet that is too simplistic. For in some sense, a proprietary service can make up *part* of the Internet when people build other (open) services that depend upon the proprietary APIs. That is part of the paradox of the 'real Internet': it is made up of independent segments that become part of the global Internet if enough other parts of the network depend upon those segments to enable the other parts to function. Some of those segments *could* be closed and proprietary. But it is impossible to know how flexible such a system would be. Engineers call such a system 'brittle', because it does not have much flex before it breaks.

As a result of all the above considerations, it is possible to discern five deployment patterns that are implicated in consolidation. It is worth considering whether each one is really a pattern that affects the Internet:

(1) **Entirely private, consolidated, API-based services that turn out to be killer applications**: These are, in effect, just new applications that live on top of the Internet. If they become socially important, then they are subject to the usual constraints on businesses and monopolies and so on. But they are not really a concern for the Internet *per se*, and we should likely respond to them without resorting to regulation of networks and of the Internet itself. That is not to say that no regulatory activity is appropriate here. In the regular humans-and-steel world, we have ways of handling straight monopolies and their relations with consumers. This might be an area for regulation. It might require novel interpretations of existing regulation, and there may be jurisdictional clashes that will have consequences *for* the Internet. But new applications *atop* the Internet are not the Internet itself.

(2) **Monopoly or consolidated API-based services that turn out to be infrastructure**: If a service becomes something on which others base additional applications or services, in some sense it becomes a part of the Internet. Indeed, building such services on top of private APIs (which are just protocols by a different name) may be an indication that an open protocol was needed in the first place. When there are multiple competing services that are built this way, it can mean that there is a policy issue to confront, because switching costs between competing services are higher than they would

be if an open protocol were available. Confronting this may require thought about incentives rather than any attempts at direct regulation.

(3) **Consolidation in standard Internet services**: In principle, anyone competent could offer standard Internet services. Consolidation in this area probably just means that a specialist commodity market is developing. In the case of a commodity, switching costs are low, so this kind of consolidation is probably not really a worry given that there are comparable substitutes (but see below). If some powerful operator starts to misbehave, another operator will almost certainly appear to provide the needed competition since the services rely on open protocols. Consolidation here might be temporarily bad, but it is very unlikely to be a stable state of affairs unless one of the other scenarios also applies.

(4) **Consolidation in standard Internet services with tied special services**: While standard services provided by specialists, described above, are a boon to consumers, they're a problem for the providers. If you offer the very same service anyone else can, differentiated only by price, your business is in a race to the bottom. Many service providers, faced with this, start to find 'extensions' of standard services that are not themselves standards. For instance, DNS providers wanted to provide DNS responses that differed depending on who asked or depending on the network-topological location of the server. The service providers usually offered this service in a proprietary manner that could not interoperate with competitors' similar offerings. These are really an example of case 2 above. An important policy question is whether it is possible to align economic incentives to turn these 'extensions' into standard cases.

(5) **Consolidation and linkage among standard and proprietary services**: When pattern 4 and patterns 1 or 2 come together, it may amount to a business model that encourages consolidation rather than one that is technologically neutral. HTTPS offers a number of advantages over certain traditional native transports, so it is often attractive to deploy a standard service over HTTPS. Yet HTTPS as a transport mechanism contains a number of features – user data available in HTTP headers, the potential for JavaScript use and openness to use or provision data at other websites – that might increase available data for a provider. This is part of the concern about DNS over HTTPS (DoH), for example: the objection is often less about the protocol itself and more about who is going to end up operating DoH servers. All of these are economic considerations strictly speaking outside the Internet architecture, but they're socially and economically important.

What all of this indicates is that, while consolidation presents serious legal and social issues, it only sometimes presents an issue for the Internet itself. It is critical that policymakers and, for that matter, all citizens recognise that the Internet itself can remain a neutral tool that provides opportunities to maximise the choices for users. But remaining that way depends, to a large extent, on the degree to which we all defend it. To maintain the Internet we want and deserve requires care in developing and choosing regulatory frameworks, responses and non-responses. I believe, with the Internet Society, that the papers contained in this issue represent a contribution to ensuring we have the Internet for everyone in the future.