On Internet Governance

[for Laura Abba]

Vinton G. Cerf
1/31/2022 1348 ET
DRAFT V 1.1

Introduction

What do we mean when we speak of “Internet Governance”?

As the third decade of the 21st Century unfolds we find ourselves in the midst of a global pandemic. Among many lessons from this continuing experience, we have found that the global Internet is playing a significant role for many of the nearly 60% of the world’s population that has access to it. Of course, the quality, cost and reliability of access varies and in every country there are areas that have little or no useful access. Nonetheless, its benefits have been apparent, especially for those able to work remotely, obtain general access to information and those relying on it for important scientific exchanges and. The latter contributed to the speed with which vaccines were developed against the SARS-COV-2 virus that leads to COVID-19 infection.

The power of the Internet and the applications that have developed using the World Wide Web platform that rides on the Internet have alerted governments, the private sector, academia and the general public to the benefits and the hazards that this system affords. The optimism of the early days of the Internet and the World Wide Web have given way to the recognition that this powerful set of technologies can and is being abused by people that do not have the general public’s best interests at heart. With the arrival of *social media* such as Facebook, Twitter, Tik-Tok and YouTube, opportunities for group formation on both global and local scales emerged. More generally, the sources of information has grown dramatically without regard to quality, thanks to these new application technologies. Opportunities for abuse of the Internet have mounted and become both more visible and potentially more harmful (think ransomware and destructive malware, misinformation and disinformation, phishing and spam email), governments have concluded that some kind of regulatory response is needed, including law enforcement.

It is in this context that this essay tries to analyze the complex challenge of introducing governance measures to preserve the utility of the Internet, the safety of citizens and to hold accountable those who abuse the privilege of access to the Internet.

Technical Perspective

It is important for parties interested in the governance of the Internet to know about the basic layering of the Internet (including the World Wide Web) because the various opportunities for governance interventions vary depending on the layer in which issues arise requiring regulation.

The Internet is designed as a layered system. The lower layers of its implementation involve the physical transport of digital information using *packet switching* technology. A useful analogy is to imagine electronic postcards that have *to* and *from* addresses and some content. Internet packets, like postcards, don't know how they are being carried. They could be transmitted on wires, coaxial cables, optical fibers, radio and satellite channels. In addition, like postcards, Internet packets don't know what information they are carrying. This ignorance is actually a major design benefit. When new transmission technologies come along, they can easily carry the digital packets of the Internet. Moreover, if a new application is developed that needs to interpret the payload of the Internet packets (think "things written on the postcard"), the network need not change because it doesn't care what the application is. While this produces a *best efforts* network, rather than one fine-tuned for specific applications, the Internet has been able to adapt to a remarkable number of applications including electronic mail, remote access to time-shared computers, streaming audio and video, video conferencing, real-time gaming, remote device control (think *Internet of Things [IOT])* and a host of others.

It is important to recognize that the basic Internet is not the World Wide Web nor is it all the applications that use it. For purposes of this essay, the Internet is a transport system that moves packets of data from source to destination. There are a number of *protocols* (think: practices, procedures, standard formats of data) that make it possible for the Internet to move packets around. Several core protocols make up the basic Internet. These are the Internet Protocol (IP), the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP). One can think of IP as a basic postcard service. IP packets have numerical addresses for all sources and destinations in the Internet and a method for finding routes from source to destination. TCP makes sure all the postcards are delivered completely and in order by introducing a layer above the IP protocol to achieve those objectives. It keeps track of re-ordering packets that might arrive out of order, re-transmitting any that might appear to be lost, filtering of duplicate packets and managing the flow of the traffic to keep the receiver from being overwhelmed. The UDP protocol gives access to low delay transport over IP without introducing any of the discipline of the TCP protocol. The result is lower latency (delay) but with the potential for disorderly arrival and receipt of duplicate traffic. The aggregate suite of protocols that make up the basic Internet are sometimes called the *TCP/IP Protocol Suite.* It includes other protocols that support routing of traffic across the global Internet, encryption of packets for confidentiality, translation of *Domain Names* (think: *cnr.it*) into iP addresses.

Applications that use the basic Internet have yet other protocols for their implementation. Remote access to time-shared computers and data centers, electronic mail transport, file transport and streaming audio and video are examples of applications that use the Internet. The

World Wide Web is implemented with its own set of protocols (*Hypertext Transport Protocol [HTTP] and Hypertext Markup Language [HTML]*) running on top of TCP/IP.

The introduction of the Apple *iPhone* in 2007 represented an inflexion point in application space with millions of applications being created for *smartphones* from multiple suppliers. The writers of these applications need only to know how the smartphone accepts and delivers data through *Application Programming Interfaces (APIs)* without knowing all the details of how mobile phones actually move voice and data over the air and through supporting fiber networks. The smartphone made the Internet more accessible and the Internet made the smartphone more useful. It is increasingly the case that smartphones are essentially endpoints in the Internet Protocol space and are directly reachable via TCP/IP and UDP protocols. The applications found on smartphones are increasingly manifestations of access to Web servers in data centers around the world. Voice communication is rapidly becoming *Voice over IP (VOIP)* as opposed to the earlier circuit-switched, analog voice of the past. This is so even on traditional telephone networks that have abandoned circuit switches for packet switches that cost less and are more flexible.

It is also vital to recognize that the Internet was designed to be expandable. It is a *Network of Networks* designed to allow an arbitrarily large number of networks to be interconnected. Each is operated independently of the others but they all follow the same standard protocols so as to achieve interoperability among all of them. Anything, anywhere on the Internet can communicate with anything else on the Internet regardless to which component networks the source and destination devices are connected.

Organizational Perspective

Just as it is important to appreciate the layered technical structure of the Internet and the World Wide Web, it is also important to understand the role of organizations in their implementation and operation. For example, some key players in the Internet environment are cable and fiber operators delivering Internet service to homes, businesses, institutions and government agencies. Similarly the wireless carriers do the same thing over the air. Others provide underlying submarine cable service linking pieces of the Internet around the world.

Other players make *routers* that implement the IP system of network interconnection or smartphones. The Internet of Things refers to a growing number of devices that are *internet-enabled* and are made by manufacturers who want to take advantage of the global connectivity of the Internet to provide devices and possibly associated services such as monitoring or software updates to the users of these devices. Still others make Domain Name resolvers and servers to translate domain names into IP addresses. Companies like Apple, Microsoft and Google make operating systems such as OSX, IOS, Windows or Android used in many smartphones, pads, laptops and IOT devices.

In addition to organizations and businesses that make Internet-enabling or Internet-enabled devices, there are *Internet service providers (ISPs)* that provide access to the Internet. These include mobile smartphone services, fiber and cable services, and low, medium and geo Earth orbit satellite services. Other organizations operate *Internet Exchange Points* that allow for efficient interconnection of many networks that make up the Internet. For the most part, these operators are focused only on the delivery of Internet packets and not on the interpretation of their payloads except to the extent that these are part of the Internet operational control system such as the *Border Gateway Protocol (BGP)* routing system.

Platform companies such as Google, Facebook, Microsoft, Amazon, IBM and others operate multiple data centers around the world, often interconnected by private fiber networks which are, in turn, connected to the public Internet. It is on these platforms that a significant fraction of Internet and Web-based services are supported. In addition, there are companies such as Akamai that operate *Content Distribution Networks* that bring content closer to users by placing it in servers that are colocated with telephone central offices and cable company head ends or nearby Internet Exchange Points.

Countless companies make application software that runs in the various data centers or CDNs to serve customers of these applications. Others create content such as Netflix and Amazon and the traditional movie studio companies which is distributed through streaming on the World Wide Web. Banks provide financial services and millions of companies create websites to serve customers. Advertising is a major component of the 21st Century Internet and drives some businesses that provide free services to users in exchange for showing them advertisements paid for by companies wishing to offer products and services to the general public.

There are, however, another class of organizations that are tied more closely to the technology of the Internet. Some of them create and maintain technical standards for the Internet among which are the Internet Engineering Task Force (IETF) sponsored by the Internet Society; the World Wide Web Consortium (W3C) for Web standards; the Organization for International Standardization (ISO), 3GPP (for mobile standards), the International Telecommunication Union (ITU) and its standards bodies ITU-R (radio) and ITU-T (telecommunication) and the Institute of Electrical and Electronics Engineers (IEEE) along with various national bodies such as the British Standards Institution (BSI), the American National Institute of Standards and Technology (NIST), the European Telecommunication Standards Institute (ETSI) and the Italian Organization for Standardization (UNI).

It is important to recognize in the context of governance that there are many organizations with an interest in the “rules of the road” for the Internet and WWW. This includes civil society and the academic community in addition to governments at all levels. It is for this reason that the notion of *Multistakeholder Internet Governance* has been a common thread in the history of the Internet and the Web. There are simply an enormous number of interested parties, parties with the power to develop and effect various norms and regulations and parties who are affected by these same rules and regulations.

Internet Governance

We come now to the crux of this essay: how is governance of the Internet to be effected with the goal of preserving the enormous value that the free flow of information has provided in the decades of operation of the Internet (since 1983) and the World Wide Web (since 1991). Multistakeholder development of policy including governments, civil society, the academic and standards communities and the private sector is highly desirable so as to include their perspectives and to understand the potential impact on various sectors of choices of policies, regulations and methods of enforcement. While enforcement of policy may fall to a smaller cohort, broadly informed policy development has been beneficial for the Internet and World Wide Web over the course of their evolution.

One important principle to urge is *subsidiarity*. That is, to apply regulatory mechanisms at the appropriate layers and to the appropriate players of the Internet ecosystem. For example, fragmenting the basic Internet in the interest of so-called *data sovereignty* threatens the free flow of information on a global scale. The value of that free flow is vital to the sharing and discovery of information that benefits everyone. If protection of information is a high priority, and it is for many things including *personally identifiable information*, one can make use of cryptographic means both to protect information in transit and at rest and to strongly authenticate users to provide only authorized access to the sensitive information. It is not necessary to confine information to specific geographies to achieve that effect. Low level geo-fragmentation of the Internet has negative consequences including the inhibition of replication of data at multiple datacenters to avoid loss even in the face of catastrophic failure.

If content is unacceptable in some contexts, blocking Internet access at the IP address level is an extremely blunt instrument. Similarly so with domain name blocking. By analogy, if someone is selling drugs illegally from an apartment, arresting all occupants of the building penalizes many innocent parties. If content is widely considered unacceptable, such as Child Sexual Abuse Material (CSAM), a potentially less blunt response is to demand its removal from serving sites. Of course, because the Internet and the Web are global, this might require the cooperation of national governments as is contemplated by the UN Secretary-General's Digital Cooperation proposals. In the absence of cooperation, governments may choose to apply much more blunt instruments.

Regulators will benefit from a refined view of the layered structure of the Internet and World Wide Web as well as an appreciation for the diversity of the ecosystem that animates them. Applying policy to the appropriate parties, recognizing their roles in the system can preserve the benefits of the Internet writ large while focusing attention and regulatory constraints more appropriately. It is worth recognizing that many of the organizations that populate the ecosystem perform multiple roles and some have vertical components that should be viewed from the layered perspective when considering regulatory responses to governance concerns.

Afterthoughts

In this brief essay, I have not attempted to deal with some really pernicious problems such as malware, denial of service attacks, the spread of misinformation and disinformation, harmful effects of social networking, the role of the general public and users in defending themselves from risk and harm. It seems valuable, however, to draw attention to some of these matters, however lightly. Users need to be literate about the potential hazards of *online life*. They need to have tools for protection such as two-factor authentication in addition to their usernames and passwords. They need to be able to detect likely *phishing* attacks. They need to avoid downloading files and software from questionable sources. They need to avoid clicking on hyperlinks without applying some critical thinking to their origins and intent.

Open Source Software has been a boon for programmers but a major hazard because careful evaluation of the software is neglected on the assumption that “all the bugs have already been found.” Sadly, this assumption often means “none of the bugs have been found.” The most recent example of this is the *Log4J* bug that allowed exploitation by running arbitrary software on the infected user’s computer. This may be the worst bug found in decades as the software is in extremely wide use. Supply chain attacks that go after critical pieces of software supplied by smaller companies not well-prepared to defend against penetration and alteration of their products have led to very serious consequences for companies and consumers. Ransomware attacks have had cascading side-effects such as the near shutdown of commerce on the US East Coast due to lack of fuel when the Colonial Pipleline operation was hit with an attack.

In our zeal to strongly identify and authenticate users, it is possible to create massive bio-data collections (faces, fingerprints, eye iris scans) that are extremely attractive to hackers interested in using or reselling this kind of information. As we recognize the potentially harmful aspects of social networking applications, companies and governments will benefit from the advice and insights from sociologists, psychologists, neuro-scientists and even anthropologists. Human behavior is complex and it is partly determined by evolutionary physiology in which emotional triggers spur reaction without much thought (e.g. fight or flight, thinking fast and thinking slow).

I have given little or no space to the whole matter of digital inclusion: making the Internet accessible (in both senses), available, affordable, reliable, safe, secure, privacy-protecting for everyone on the planet. There are many barriers to achieving global access and these, too, deserve governance attention and have not been given adequate treatment in this short essay.

There is no question in my mind that we need to attend to the potential risks we face with computer-based, networked services, but I believe strongly that we must be mindful of losing much of the already demonstrated benefits if we are not cognizant of the complexity of the ecosystem we have created, its structure and players in it.