



INTERNET GOVERNANCE PAPERS

PAPER NO. 2 — AUGUST 2013

Internet Points of Control as Global Governance

Laura DeNardis



INTERNET GOVERNANCE PAPERS

PAPER NO. 2 — AUGUST 2013

Internet Points of Control as Global Governance

Laura DeNardis

Copyright © 2013 by The Centre for International Governance Innovation.

The opinions expressed in this publication are those of the author and do not necessarily reflect the views of The Centre for International Governance Innovation or its Operating Board of Directors or International Board of Governors.



This work was carried out with the support of The Centre for International Governance Innovation (CIGI), Waterloo, Ontario, Canada (www.cigionline.org). This work is licensed under a Creative Commons Attribution — Non-commercial — No Derivatives License. To view this license, visit (www.creativecommons.org/licenses/by-nc-nd/3.0/). For re-use or distribution, please include this copyright notice.

Cover and page design by Steve Cross.

ACKNOWLEDGEMENT

CIGI gratefully acknowledges the support of the Copyright Collective of Canada.



CONTENTS

About the Author 1

About Organized Chaos: Reimagining the Internet Project 2

Acronyms 2

Executive Summary 3

Introduction 3

Global Struggles Over Control of CIRs 5

Governance via Internet Technical Standards 8

Routing and Interconnection Governance 10

Emerging International Governance Themes 12

Works Cited 14

About CIGI 15

ABOUT THE AUTHOR

Laura DeNardis

Laura DeNardis, CIGI senior fellow, is an Internet governance scholar and professor in the School of Communication at American University in Washington, DC. Her books include *The Global War for Internet Governance* (forthcoming 2014), *Opening Standards: The Global Politics of Interoperability* (2011), *Protocol Politics: The Globalization of Internet Governance* (2009) and *Information Technology in Theory* (2007, with Pelin Aksoy). She served as the executive director of the Information Society Project at Yale Law School from 2008–2011, and is a co-founder and co-series editor of the MIT Press Information Society book series. She currently serves as the elected vice-chair of the Global Internet Governance Academic Network. Laura holds an A.B. in engineering science from Dartmouth College, an M.Eng. from Cornell University, a Ph.D. in science and technology studies from Virginia Tech, and was awarded a postdoctoral fellowship from Yale Law School.

ABOUT ORGANIZED CHAOS: REIMAGINING THE INTERNET PROJECT

Historically, Internet governance has been accomplished *en passant*. It has emerged largely from the actions of computer scientists and engineers, in interaction with domestic legal and regulatory systems. Beginning at least with the 2003–2005 World Summit on the Information Society process, however, there has been an explicit rule-making agenda at the international level. This strategic agenda is increasingly driven by a coalition of states — including Russia, China and the Arab states — that is organized and has a clear, more state-controlled and monetary vision for the Internet. Advanced industrial democracies and other states committed to existing multi-stakeholder mechanisms have a different view — they regard Internet governance as important, but generally lack coherent strategies for Internet governance — especially at the international level. Given the Internet’s constant evolution and its economic, political and social importance as a public good, this situation is clearly untenable.

A coherent strategy is needed to ensure that difficult trade-offs between competing interests, as well as between distinct public values, are managed in a consistent, transparent and accountable manner that accurately reflects public priorities. Guided by these considerations, CIGI researchers believe they can play a constructive role in creating a strategy for states committed to multi-stakeholder models of Internet governance.

In aiming to develop this strategy, the project members will consider what kind of Internet the world wants in 2020, and will lay the analytical groundwork for future Internet governance discussions, most notably the upcoming decennial review of the World Summit on the Information Society. This project was launched in 2012. The Internet Governance Paper series will result in the publication of a book in early 2014.

ACRONYMS

ASN	Autonomous System Number
BGP	Border Gateway Protocol
CIRs	critical Internet resources
DNS	Domain Name System
gTLD	generic top-level domain
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IETF	Internet Engineering Task Force
IGF	Internet Governance Forum
IP	Internet Protocol
IPv4	IP version 4
ISOC	Internet Society
ITU	International Telecommunication Union
IXP	Internet Exchange Point
LINX	London Internet Exchange
NTIA	National Telecommunications and Information Administration
RFCs	Request for Comments
RIRs	Regional Internet Registries
TLD	top-level domain
W3C	World Wide Web Consortium
Wi-Fi	Wireless Fidelity

EXECUTIVE SUMMARY

The distributed nature of Internet infrastructure and relatively malleable user engagement with content can misleadingly create the impression that the Internet is not governed. When Internet governance does rise to media or public prominence, this usually involves high-profile controversies such as the Egyptian government cutting off citizen Internet access or government-delegated censorship requests for Google to delete politically sensitive content. These are examples of Internet content governance via infrastructure. But beneath this layer of content, at much more technologically concealed layers, coordinated and sometimes centralized governance of the Internet's technical architecture is necessary to keep the network operational, secure and universally accessible. This governance is enacted not necessarily through traditional nation-state authority but via the design of technical architecture, the policies enacted by private industry and administration by new global institutions. While these coordinating functions perform highly specialized technical tasks, they also have significant economic and political implications.

This paper explains how the Internet's core technical architecture is governed and how global public policy decisions are co-produced within this governance framework. It describes three fundamental control functions necessary for the Internet to operate: control of critical Internet resources (CIRs), such as names and numerical addresses; governance via Internet standards; and governance of routing and interconnection. These core areas of technical coordination collectively enable the defining Internet characteristics of interoperability, accessibility and universality, but also shape public policy related to information access, individual rights, security, innovation and economic competition. It concludes by highlighting common themes among these areas, including the privatization of governance and values


tensions mediated at Internet control points, and by raising several open governance issues, including proposed changes in interconnection agreements and architectural changes agonistic to universal interoperability.

INTRODUCTION¹

Internet governance is a capacious topic, involving the administration of the Internet's technical architecture and the formulation of public policy around this architecture. Key questions include *what* infrastructural Internet components are currently governed and *why*; *how* coordination and control occurs and *by whom*; and what are the broader policy implications of this coordination.

Internet governance can be broken down into a number of different taxonomies, but one way to divide its tasks is into the following six broad areas (see Figure 1): architecture-based intellectual property rights enforcement; the policies enacted by information intermediaries; cyber security governance; governance of routing and interconnection; Internet standards governance; and control of CIRs.

¹ A more extensive treatment of Internet governance is presented in Laura DeNardis (forthcoming 2014), *The Global War for Internet Governance*, Yale University Press.

Figure 1: Core Tasks of Internet Governance

The Internet is comprised of independently operated networks and countless types of hardware, software and standards, but the common denominator technology defining when someone is “on the Internet” is the use of the Internet Protocol (IP). This core common architecture contributes to the Internet’s basic characteristics of universality, interoperability and accessibility, and can be used as an organizing principle for understanding fundamental technical mechanisms of Internet governance. In the context of twenty-first century Internet usage, many general characteristics, including the ability to access the universal Internet from almost anywhere in the world, can be taken for granted. As a contribution to the United Nations Internet Governance Forum (IGF), a multi-stakeholder coalition has been examining the questions of what are the Internet’s architectural principles, what are its core values and how are these values being upheld or diminished as the Internet evolves?² With the exception of repressive political contexts of censorship, the Internet’s core values are universality, interoperability and accessibility. Someone in Sydney, Australia, can generally access or upload the same information as

someone in Toronto, Canada. Most points anywhere on the Internet can reach any other point. It is a single, universal network. These characteristics are not a given, but are designed into the Internet’s technical architecture.

This paper focusses on the functions of Internet governance most closely related to enabling the Internet’s universality, accessibility and interoperability, and explains how global public policy decisions are co-produced within this governance. Specifically, it describes three fundamental control functions necessary for the Internet to operate: control of CIRs; governance via Internet standards; and governance of routing and interconnection. Coordination of unique virtual resources, such as binary Internet addresses and alphanumeric domain names, is necessary for the Internet to operate. This paper describes the system of control over the distribution of these resources — including by new global institutions such as the Internet Corporation for Assigned Names and Numbers (ICANN) and Regional Internet Registries (RIRs) — and explains policy issues around the administration of these resources. Internet standards are the rules, or protocols, that computing devices follow to ensure interoperability with other computing devices that also adhere to these standards. These rules are set by many standards-setting institutions, including

² See, for example, the transcript from the third meeting of the Dynamic Coalition on Core Internet Values from the IGF (2012), Baku, Azerbaijan, November 8, available at: <http://wsms1.intgovforum.org/2012/Meetings/dynamic-coalition-core-internet-values>.

the Internet Engineering Task Force (IETF) and the World Wide Web Consortium (W3C). Although they perform a range of very specific technical functions related to interoperability, they also shape public policy related to information access, individual rights and security. Internet interconnection agreements usually involve private contractual arrangements among network operators to connect bilaterally or at shared Internet Exchange Points (IXPs), raising governance questions about who can connect and under what economic terms.

This paper concludes by highlighting common themes among these technical areas, including the privatization of governance and the increasing phenomenon of Internet control points serving as sites of mediation for conflicting global values. It also raises open issues of Internet governance, including proposed changes in how interconnection agreements work and architectural changes that could diminish universal interoperability.

GLOBAL STRUGGLES OVER CONTROL OF CIRs

Just as basic functioning in the offline world requires scarce natural resources, such as water and energy, the Internet's basic functioning requires finite virtual resources. These CIRs are the unique binary and alphanumeric identifiers that comprise the Internet's system of naming and addressing, and the massive, distributed Domain Name System (DNS) that translates between the names that people use to access an online site and the binary addresses computers use to locate and route information to that site.

Global tensions over the control and distribution of these resources, as well as the enactment of policy issues around these resources, have been a long-standing struggle of Internet governance. Conflicts

have often centred on the historic relationship the US government has had with certain control aspects of these resources. Newer narratives have expressed alarm over a possible "takeover" by the United Nations and, in particular, its specialized subagency for information and communication technology, the International Telecommunication Union (ITU). Still other governance debates have questioned the global coordinating role of new global institutions like ICANN. This section identifies the technical resources at the centre of these struggles, explains the institutional system of governance that currently oversees these resources and presents several pressing international public policy concerns over how governance over CIRs does or should occur.

IP addresses are the unique binary numbers every device using the Internet possesses, either permanently or assigned temporarily for a session. The format of Internet addresses is specified by the IP standard. The long-standing version of IP, known as IP version 4 (IPv4), assigns 32 bits (32 zeros and ones) to each binary address — for example, 00010011001010001000000100100001. Internet users might not be directly cognizant of a specific binary address, although more customarily they might have seen a number such as 19.40.129.33, which is the dotted decimal notation for the above binary address. IP addresses are at the heart of how the Internet routing functions, because they are used by routers to transmit information to its destination over the most expeditious path.

This design feature of a 32-bit address mathematically produces a pool of 2^{32} , or roughly 4.3 billion unique Internet addresses, an insufficient number to meet the demands of global Internet growth. A newer standard, IPv6, expands the address length to 128 bits, providing a unique pool of 2^{128} , or 340 undecillion addresses. The new standard has long been available and implemented in products, but

for a variety of reasons has not been deployed, most notably because the new standard is not backward compatible with IPv4.

Networks interconnecting to form the global Internet also each possess a unique binary number called an Autonomous System Number (ASN). The assignment of a unique ASN is necessary for network operators to exchange information and is a prerequisite for becoming a network operator. How this system of interconnection among ASNs works is explained in a later section.

Fortunately, users accessing an online site do not have to enter a binary address. They enter an alphanumeric domain name (for example, www.cigionline.org). Like IP addresses, each domain name must be globally unique. The DNS is the universal technology of Internet governance that translates between the domain names that humans use and the binary addresses computers use. The DNS is a massive, distributed database management system stored on servers around the world and performing this address resolution function for billions upon billions of transactions per day.

The Internet can only operate with a unique name and number space (Internet Architecture Board, 2000). This technical design decision to use globally unique and finite stores of name and number identifiers has produced requirements for a specific kind of governance. Someone has to centrally coordinate the allocation of IP address blocks to ensure that each assigned number is globally unique. Someone decides how many of these become delegated to regions and what regional institutions can be given the authority over how to locally distribute numbers and on what basis. Someone determines what institutions can be assigned ASNs in order to become network operators. Someone has to assign globally unique domain names to end users, authorize the introduction of new top-level domains

(TLDs) (such as .com or .books) and adjudicate domain name trademark disputes that can arise. Someone is responsible for keeping the definitive record of how to resolve names into numbers for each TLD and for the root zone file containing the master, most centralized list that dictates how each TLD maps to binary addresses.

ICANN has centralized authority over most governance functions related to domain names and addresses, although this authority is further delegated to a complex and distributed mosaic of other, mostly private organizations. This governance framework is now a stew of acronyms, but at one point in Internet history, a single individual, the late Jon Postel, distributed and tracked Internet numbers and various unique identifiers that kept the Internet operational. The function he and his colleagues performed decades ago was, and still is, called the Internet Assigned Numbers Authority (IANA), although now a function under ICANN, a private, non-profit corporation that the US government contracted in 1998 to coordinate names and numbers and administer the Internet's root servers.³

IANA allocates addresses for regional assignment to five RIRs: AfriNIC — the African Network Information Centre; APNIC — the Asia Pacific Network Information Centre; ARIN — the American Registry for Internet Numbers (Canada, United States, North Atlantic islands); LACNIC — the Latin America and Caribbean Network Information Centre; and RIPE NCC — Réseaux IP Européens Network Coordination Centre (Europe, Middle East, parts of central Asia). These member-funded institutions are quite powerful because they control the allocation of Internet addresses in their respective regions. The

3 For a lengthy history of the evolution of root management, see Milton Mueller (2002), *Ruling the Root: Internet Governance and the Taming of Cyberspace*, MIT Press.

RIR function is an area of privatized governance that is neither under government control nor market based.

In the generic top-level domain (gTLD) name space, ICANN also accredits the hundreds of “registrars” (for example, Go Daddy) that sell domain name registrations to institutional and individual customers, and also delegates authority to the registry operators responsible for maintaining and distributing the authoritative mapping of names and associated IP addresses for every domain name contained within a TLD. When one considers the number of registrars, registry operators and RIRs that allocate IP addresses, there is an enormous number of institutions that provide coordinating governance functions over CIRs.

ICANN has become quite internationalized in composition and structure over the years, but the historic relationship between the US Department of Commerce and the CIR governance framework has remained a contentious governance question. For example, the IANA function under ICANN is specifically authorized by a contract with the National Telecommunications and Information Administration (NTIA) of the US Commerce Department (NTIA, 2012). Jurisdictional authority over the root zone file also resides with the NTIA, although delegated to IANA and to a private US company called VeriSign (US Department of Commerce, 2013). The United States’ delegated control of the root zone file, its contract with the IANA and its historic relationship with ICANN have placed the question of US control at the centre of global power struggles over the Internet, often based on principle rather than substantive policy concerns related to how the root zone file is actually administered.

Many tangible policy concerns arise more broadly in the area of CIR governance. One policy area with

considerable implications to freedom of expression, innovation and property rights is the massive expansion of TLDs. The number of TLDs has risen gradually over the years, beginning with gTLDs (such as .com, .org and .edu) to the introduction of country code TLDs (such as .ca, .uk and .in). Domain names originally were relegated to Latin alphabet characters, but have also expanded to scripts that enable native languages, including Arabic, Cyrillic, Chinese characters and other scripts. A more recent controversy has been ICANN’s strategy to dramatically expand the number of available TLDs.

In response to a 2012 call for proposals for new gTLDs, ICANN received nearly 2,000 proposals ranging from .blog, .shop, .apple to .books.⁴ Those companies proposing a new TLD paid US\$185,000 for the application alone. Applicants would also commit to being responsible for the registry, raising concerns about whether there would be a free market for any entity wanting to register, for example, a .cloud domain name, or whether there would be anti-competitive behaviour around new gTLDs. On one hand, the expansion simply parallels the growth of the Internet and increases spaces for innovation and expression. On the other hand, this expansion makes it more challenging for trademark holders to protect their intellectual property rights and could expand spaces for media piracy. Either way, the expansion of domain names will engender new Internet governance struggles. Such tensions quickly emerged after the application process for new gTLDs, such as conflicts between trademarked company names and geographical regions. For example, the companies Patagonia and

4 For the list of applicants for new gTLDs published on ICANN’s website, see ICANN (2012), “Reveal Day 13 June 2012 – New gTLD Applied-For Strings,” June 13, 2012, available at: <http://newgtlds.icann.org/en/program-status/application-results/strings-1200utc-13jun12-en>.

Amazon applied for gTLDs of their respective names (.patagonia and .amazon), but countries with the Amazon and Patagonia regions within their borders objected to these applications.

GOVERNANCE VIA INTERNET TECHNICAL STANDARDS

Human languages and social conventions provide the rules that govern basic human functions related to communication and interactions. Obvious examples include regulations related to transportation and driving (such as which side of the road to drive on), cultural conventions for greeting someone and standards for how to address a letter. These standards are social constructs that can vary by culture. Just as humans adhere to standards that enable the exchange of information, so it is with digital devices. These standards, also called protocols, are the Internet's common language, the specifications that establish universal formats for how to digitally encode information, how to address the information so that it can reach its destination, or how to compress, encrypt or otherwise manipulate binary code so that it can be interoperably exchanged among any device connected to the Internet. Standards are neither software code nor hardware; they are written specifications dictating how to develop software and hardware to be compatible with any other type of software and hardware that also adheres to these specifications.

Prior to the development and adoption of the Internet's core family of protocols, known as Transmission Control Protocol/Internet Protocol (TCP/IP), devices made by one company (such as IBM) could not exchange information with devices made by another developer (such as Apple). Networks of computers, as well as online services such as American Online or CompuServe, relied on proprietary, closed protocols inaccessible for other companies to access

or implement in products. The Internet's universality and interoperability is made possible because hardware and software manufacturers now use common technical standards. The open publication and availability of these standards has contributed to the Internet's rapid progression of innovation and growth and has been described as the "most formidable regulatory regime that has governed the Internet to date" (Weiser, 2001).

Most people have heard of some of the standards they use for the everyday exchange of digital information, although it might not register that these household names are actually standards. Examples include the Wireless Fidelity (Wi-Fi) family of standards for wireless local access or formats for digitizing media, particularly the MP3 format for digitally encoding and compressing audio. Accessing a website from a browser relies upon the Hypertext Transfer Protocol and making a voice call over the Internet relies on Voice over Internet Protocol. The vast majority of the thousands of protocols necessary for exchanging information over the Internet are not visible to users, yet they are necessary for creating order to binary streams of information and for ensuring interoperability among devices adhering to these standards.

While technical standards perform quite esoteric technical functions, they also produce global economic and political effects and, to a certain extent, enact public policy in areas that are traditionally carried out by governments (DeNardis, 2009). Economic analyses of technical standardization tend to focus on the salutary network effects of a standard on innovation, market efficiency, global trade and national economic competitiveness.⁵ Internet standards provide the blueprints for entrepreneurs

5 See, for example, Knut Blind (2004), *The Economics of Standards: Theory, Evidence, Policy*, Edward Elgar Cheltenham.

to use to create new product innovations. Use of an industry standard minimizes risk because the manufacturer is assured that the product will “work” in the marketplace. In this regard, the availability of open Internet standards, particularly those without restrictions on their implementation, promotes conditions for innovation and a free market of multiple competing, but compatible, products. As in many other sectors, standardization in the information and communication technology sector facilitates international trade and export competitiveness.

Beyond these economic effects, standards shape public policy in more direct political ways. They are the infrastructural foundations for global trade and the digital public sphere, but their design and constitution create public policy in areas as politically charged as privacy, accessibility and other individual civil liberties. Encryption standards and, in particular, the strength or key length of encryption standards, mediate between conflicting social values of the expectation for privacy and responsibilities of law enforcement and national security. The engineering design of IP established the requirement for a globally unique virtual identifier for each exchange of information over the Internet, a characteristic with its own privacy implications. The “DoNotTrack” protocol seeks to provide a privacy option for those wishing to circumvent the increasing forms of behavioural tracking fuelling online advertising. Web accessibility standards design accessible features into technical specifications that address differences in vision, speech, movement, cognition and hearing — such as enabling closed captioning in online video. The W3C (2008) has a Web Accessibility Initiative seeking to make the Internet more accessible for individuals with various physical or cognitive impairments.

Standards can also be political, in the sense of being associated with controversy. The protocol BitTorrent serves as a prime example. From an engineering standpoint, it performs an efficient technique for transmitting large files over the Internet by breaking files into fragments and storing these fragments on distributed end computers. But from an economic and political perspective, and based on how the protocol is used, it is almost universally associated with piracy of digital media.

The policy implications of Internet standards raise the obvious governance question of how these standards are procedurally established and by whom. Dozens of distinct standards-setting organizations establish standards used in the global Internet. The IETF is one of the primary institutions setting the core, universal protocols necessary for the Internet to operate. As previously stated, if there could be a simple and precise technical infrastructure definition of the Internet, it would be the ability to reach another device via IP. The IETF has developed IP and other core networking standards for the Internet. It was formally founded in 1986, but is a direct derivative of the core Internet engineering community tracing back to the formative 1970s. More recently, the IETF was placed under an umbrella organization known as the Internet Society (ISOC), a non-profit member organization formed in 1992 and tasked with keeping the Internet operational, open and transparent (ISOC, 2013). IETF standards, along with other documents describing procedures and technical information, are electronically housed in an archive known as the Request for Comments (RFCs) series, which are records of technical specifications that date back to 1969.⁶

6 All of the Internet RFCs are freely accessible on the IETF website, available at www.ietf.org.

The W3C establishes the bulk of standards for the Web. Web inventor Tim Berners-Lee founded the W3C in 1994 to promote standardization efforts ensuring interoperability among what were then emerging and often competing Web products developed by different companies. W3C standards, called “Recommendations,” have included critical interoperability specifications, like Hypertext Markup Language and Extensible Markup Language, that explain how to encode Web information in formats that can be interpreted by any browser.⁷

Important standards work is also done by the Institute of Electrical and Electronics Engineers, which sets Ethernet LAN standards and Wi-Fi specifications, and the ITU, which has historically provided telecommunication-related standards in areas such as Internet telephony. Some standards organizations are official national bodies, such as the Standardization Administration of China or the American National Standards Institute. These are just a few examples of the standards bodies that collectively design the rules for how digital devices exchange information.

Given the significant public policy implications of standards, the procedures by which they are established is an important question related to legitimacy. The IETF is a fairly open organization, exhibiting democratic principles of openness both procedurally and in terms of how an IETF standard can be implemented. The open-participation norms of the IETF allow anyone to contribute to design efforts and decisions are based on “rough consensus and working code.” IETF standards can be freely and transparently accessed by anyone, providing an avenue for public accountability and oversight, and promoting innovation by allowing manufacturers

to develop products based on the standard. The IETF has also traditionally given preference to open standards allowing any manufacturer to develop products based on a standard with minimal or no intellectual property restrictions on its use. The W3C is a similarly open standards organization, but others have varying procedures when it comes to who can participate, the degree of procedural transparency, whether the standard is published for other to use to innovate and the degree of intellectual property restrictions on the use of the standard.

ROUTING AND INTERCONNECTION GOVERNANCE

A significant global Internet governance concern related to keeping the Internet accessible and universal is the question of how independent networks conjoin to form the global Internet. The Internet is a collection of independent systems operated by mostly private companies that interconnect to create a universal Internet. Some of these networks are large telecommunication providers, such as AT&T, Bell Canada and Korea Telecom. Enormous content companies such as Google and Facebook also operate their own networks. Other types of networks are a class of providers known as content delivery networks that content companies hire to efficiently distribute (and replicate and load balance) content on servers located around the world.

The companies that run these networks collectively form the universal Internet because they agree to adopt a set of common standards that enable interoperability among their networks, and to physically and logically connect at interconnection points. They also make private economic agreements to handle and forward traffic originating or terminating on their respective networks. This section addresses several governance questions

⁷ All W3C standards and drafts can be accessed online at www.w3.org/TR/.

related to core mechanisms of interconnectivity: how does routing work among heterogeneous networks and is this process adequately stable and secure; what private economic arrangements do network operators make to interconnect; and what role do IXPs play in interconnection governance and what are the global policy implications of IXP distribution? The final section of this paper also raises a prospective policy issue related to proposals to restructure, possibly through government interventions, the economic and institutional mechanisms of this interconnection.

Understanding interconnection begins with understanding routing; understanding routing requires understanding “autonomous systems.” Autonomous systems are routing domains, or collections of routers. Each autonomous system possesses a unique binary number known as an ASN. It also announces to the rest of the Internet a consistent routing policy and manages a set of IP addresses that can be reached within or through the system. The Internet’s routing infrastructure is its central circulatory system and the foundation of how information is transmitted from point A to point B. The Internet is an enormous packet-switching network in that when information is transmitted, it is divided into smaller segments called packets and transmitted via routers over the most expeditious path to its destination. Each packet is comprised of payload (the actual content of the information), along with accompanying administrative information such as the binary address indicating the information’s destination location. Routing algorithms and tables help routers optimize routes and minimize the latency, or delay, in transmitting information.

Routing within an autonomous system uses an interior gateway protocol, which helps each router make decisions about where to next direct a packet. Routing between autonomous systems uses an

exterior routing protocol called Border Gateway Protocol (BGP), which is one of the most important technical protocols providing the Internet’s universality and interoperability because it executes the exchange of information among networks. It dictates how networks should announce reachability, or the routes each autonomous system can reach.

This system of interconnection is, to a certain extent, based on trust among network operators. Networks assume that the routes advertised via BGP by a neighbouring system are globally accurate. In an infamous example of the fragility of this system, YouTube became temporarily unavailable in 2008. In an effort to comply with a Pakistan government order to ban access to YouTube in the country, Pakistan Telecom redirected the collection of IP addresses associated with YouTube into a digital void. But instead of relegating this redirection to local routers, the company also advertised these redirected routes outwardly, which was in turn replicated across the Internet and caused the temporary blockage. The Internet’s interconnection system has historically been fairly stable, but it does have this inherent security vulnerability in routing infrastructures. The Internet engineering community has been working on a system to secure this process via public key encryption, similar to the certificate authority system for cryptographically authenticating websites.⁸

In terms of physically interconnecting, networks conjoin bilaterally in a network operator’s premises or at large, shared IXPs. The shared connection points are a crucial part of the Internet’s physical

8 For example, Resource Public Key Infrastructure is the evolving encryption system being developed in an IETF working group called Secure Inter-Domain Routing. This system would assign to each network a digital certificate that authenticates that a network has the authority to announce the collection of Internet addresses under its purview.

infrastructure and probably the least “cloud-like” part of the network, involving buildings, cables, banks of network switches, and other physical and logical infrastructure. In the early history of the Internet, the first four interconnection points were all located in the United States and governmentally facilitated. The privatization of shared interconnection points began in the early 1990s and has grown to hundreds of locations around the globe.

To convey some scale of large IXPs, those located in large cities are often distributed across numerous distinct data centres/buildings and connected by a fibre optic metropolitan area network. Some, such as the London Internet Exchange (LINX) and the Duetscher Commercial Internet Exchange, connect hundreds of distinct network operators or content companies such as China Telecom, British Telecommunications, Facebook, Google and Akamai.⁹ These private operators pay a membership fee to connect to the IXP and also make separate agreements about the nature of how they will handle each other’s traffic.

IXPs are crucial not only to the function of the global Internet, but also to the economic and political autonomy of nations and regions. They serve as information gateways to the rest of the world and they promote the economically and technically efficient exchange of traffic directly between a nation’s network operators, rather than relying on switching facilities located in another country.

Apart from how networks physically and virtually conjoin is the issue of the financial arrangements they make for this interconnection. Historically, these arrangements have been private contractual agreements to either engage in settlement-free or paid interconnection. Although this is an

oversimplification of the many ways these private agreements are made, settlement-free peering is generally an agreement network operators make to exchange traffic without any financial obligation to each other. In other cases, peering occurs but involves an asymmetrical paid arrangement whereby one operator pays the other for this mutual peering. Other forms of paid interconnection include transit arrangements in which (usually) a smaller operator pays a larger network operator for transit connection to the global Internet.

The market for these peering arrangements is not as much based on optimizing technical or economic efficiency across the collective Internet, but on which network operators have incumbent market advantage. Not surprisingly, there have been calls to regulate and change the commercial nature of Internet interconnection, a topic raised in the next section.

EMERGING INTERNATIONAL GOVERNANCE THEMES

Control over CIRs, protocols and interconnection are the three areas of Internet governance most closely associated with preserving the Internet’s core characteristics of universality and interoperability. Examining these infrastructural-based policy areas helps elucidate several common themes about international governance. First, the Internet is already governed. As Mark Raymond and Gordon Smith (2013) explain, “The Internet has never been an ungoverned space. Even in its earliest days, it had ‘rules of the road.’ In fact, if not for such rules, the Internet would not — could not — exist.” This governance is hybridized, multi-stakeholder and highly privatized. National governments oversee some aspects of Internet policy (for example, computer fraud and abuse, antitrust, privacy) within their borders or globally via international treaties,

⁹ The entire list of full members of the LINX is available online at www.linx.net/pubtools/member-techlist.html.

but the policies enacted in deep levels of global infrastructure transcend these boundaries. The design of technical architecture through standards setting and configurations of infrastructure are making public interest decisions in areas as diverse as individual freedom and global innovation policy. Global institutions of Internet governance — whether standards organizations, RIRs, registry operators or the ICANN establishment — are private companies, informally organized international consortia primarily comprised of private industry representatives or incorporated not-for-profit corporations.

Ongoing Internet governance concerns have centred on the question of procedures for maximizing the legitimacy of these institutions, as well as determining the role of traditional governments in overseeing critical Internet infrastructure and mediating the substantive public policy issues instantiated in this infrastructure. In Internet governance debates, there has been a great interest in “preserving traditional multi-stakeholder governance” that seeks to balance governmental power, private industry self-interest, traditional economic markets and civil society. Finding this appropriate balance of powers is a context-dependent and technically and institutionally complex question. In some areas, it may be appropriate to have no governmental involvement, while other areas fall within the traditional jurisdictional bounds of democratic governance.

As history and present circumstances indicate, how Internet governance works is hardly static but in constant flux. The following two examples of prospective shifts in Internet governance are both related to the infrastructure areas explained in this paper, and both would have potentially significant implications for the Internet’s universality. One prospective change relates to international interest

in regulating, or at least facilitating, Internet interconnection. Since the commercialization of the Internet, interconnection among network operators has primarily involved private agreements. There is a long history of calls for direct governmental regulation of this interconnection, particularly over concerns about promoting greater interconnection in emerging markets, preventing anti-competitive practices and promoting “fair” compensation arrangements. A proposal related to government involvement in interconnection, although ultimately not advancing, emerged prior to an international UN conference to discuss revisions to the international treaty known as the International Telecommunication Regulations (European Telecommunications Network Operators’ Association, 2012). Imposing a telecommunication interconnection payment model on Internet interconnection, presumably enforced by governments, would be a significant departure from how interconnection has organically and rapidly grown, and could present a range of unintended or intended consequences — such as creating new concentrated points for government censorship, surveillance and politically motivated interconnection blockages, or creating economic disincentives for major content companies to interconnect in countries seeking to levy a payment, or tax, on content companies. This could potentially fragment the Internet based on political manipulation or on where content companies would economically agree to have their content accessed.

Another related change to Internet governance norms and universality would involve an increasing turn to the DNS for content control. It is already used as a mechanism for censorship in repressive information contexts, as well as being used in the United States for intellectual property rights enforcement whereby an authoritative Internet registry redirects the Internet address resolution associated with a site infringing trademark or

copyright. This only jurisdictionally works when the Internet registry controlling a domain's address resolution resides under the jurisdiction of the country in which the government requests domain name blocking. One aspect of the controversy over the SOPA/PIPA bills in the United States was that it would have expanded the government's ability to block domain names by asking local Internet service providers to alter the address resolution records they receive from authoritative registry operators located abroad. This practice would be controversial because it would fragment the Internet's universality depending on country and possibly create security and stability challenges to the DNS.

These examples further suggest how technologies of Internet governance have become the new global spaces mediating conflicting values, such as requirements for intellectual property rights enforcement and norms of Internet security and freedom. Having the attention of international governance experts and the public alike will be vital as decisions about the future of Internet governance unfold over the next decade.

WORKS CITED

- DeNardis, Laura (2009). "Political and Economic Implications of Protocols: A Framework," in *Protocol Politics: The Globalization of Internet Governance*. The MIT Press.
- European Telecommunication Network Operators' Association (2012). "CWG-WCIT12 Contribution 109." Council Working Group preparation for the 2012 World Conference on International Telecommunications, June 6.
- Internet Architecture Board (2000). "IAB Technical Comment on the Unique DNS Root," May.
- ISOC (2013). Mission Statement. Available at: www.internetsociety.org/who-we-are.
- NTIA (2012). "Commerce Department Awards Contract for Management of Key Internet Functions to ICANN," press release, July 2. Available at: www.ntia.doc.gov/press-release/2012/commerce-department-awards-contract-management-key-internet-functions-icann.
- Raymond, Mark and Gordon Smith (2013). "Reimagining the Internet: The Need for a High-level Strategic Vision for Internet Governance, 2015–2020." CIGI Internet Governance Paper Series No. 1.
- US Department of Commerce (2013). "Cooperative Agreement No. NCR 92-18742." Available at: www.ntia.doc.gov/page/verisign-cooperative-agreement.
- W3C (2008). "Web Content Accessibility Guidelines," WCAG 2.0, December 11. Available at: www.w3.org/TR/WCAG20/.
- Weiser, Philip J. (2001). "Internet Governance, Standards Setting, and Self-Regulation," *Northern Kentucky Law Review* 28, no. 822.

ABOUT CIGI

The Centre for International Governance Innovation is an independent, non-partisan think tank on international governance. Led by experienced practitioners and distinguished academics, CIGI supports research, forms networks, advances policy debate and generates ideas for multilateral governance improvements. Conducting an active agenda of research, events and publications, CIGI's interdisciplinary work includes collaboration with policy, business and academic communities around the world.

CIGI's current research programs focus on four themes: the global economy; global security; the environment and energy; and global development.

CIGI was founded in 2001 by Jim Balsillie, then co-CEO of Research In Motion (BlackBerry), and collaborates with and gratefully acknowledges support from a number of strategic partners, in particular the Government of Canada and the Government of Ontario.

Le CIGI a été fondé en 2001 par Jim Balsillie, qui était alors co-chef de la direction de Research In Motion (BlackBerry). Il collabore avec de nombreux partenaires stratégiques et exprime sa reconnaissance du soutien reçu de ceux-ci, notamment de l'appui reçu du gouvernement du Canada et de celui du gouvernement de l'Ontario.

CIGI MASTHEAD

Managing Editor, Publications

Carol Bonnett

Publications Editor

Jennifer Goyder

Publications Editor

Sonya Zikic

Assistant Publications Editor

Vivian Moser

Media Designer

Steve Cross

EXECUTIVE

President

Rohinton Medhora

Vice President of Programs

David Dewitt

Vice President of Public Affairs

Fred Kuntz

Vice President of Finance

Mark Menard

COMMUNICATIONS

Communications Specialist

Kevin Dias

kdias@cigionline.org

1 519 885 2444 x 7238



57 Erb Street West
Waterloo, Ontario N2L 6C2, Canada
tel +1 519 885 2444 fax +1 519 885 5450
www.cigionline.org