




**CHATHAM
HOUSE**
The Royal Institute of
International Affairs

Research Volume One

Global Commission on Internet Governance

A Universal Internet in a Bordered World

Research on Fragmentation, Openness and Interoperability



Research Volume One

Global Commission on Internet Governance

A Universal Internet in a Bordered World

Research on Fragmentation, Openness and Interoperability



**CHATHAM
HOUSE**
The Royal Institute of
International Affairs

Published by the Centre for International Governance Innovation and the Royal Institute of International Affairs

The copyright in respect of each chapter is noted at the beginning of each chapter.

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

This work was carried out with the aid of a grant from the International Development Research Centre (IDRC), Ottawa, Canada.

The views expressed herein do not necessarily represent those of IDRC or its Board of Governors.



This work is licensed under a Creative Commons Attribution — Non-commercial — No Derivatives License. To view this licence, visit (www.creativecommons.org/licenses/by-nc-nd/3.0/). For re-use or distribution, please include this copyright notice.

Centre for International Governance Innovation, CIGI and the CIGI globe are registered trademarks.



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

**CHATHAM
HOUSE**
The Royal Institute of
International Affairs

10 St James's Square
London, England SW1Y 4LE
United Kingdom
tel +44 (0)20 7957 5700 fax +44 (0)20 7957 5710
www.chathamhouse.org

TABLE OF CONTENTS

About the Global Commission on Internet Governance	v
Preface	vi
<i>Carl Bildt</i>	
Section One: A Universal and Open Internet	7
Introduction: One Internet: An Evidentiary Basis for Policy Making on Internet Universality and Fragmentation.	8
<i>Laura DeNardis</i>	
Chapter One: On the Nature of the Internet	21
<i>Leslie Daigle</i>	
Section Two: The Economics of Openness and Fragmentation.	40
Chapter Two: Addressing the Impact of Data Location Regulation in Financial Services	41
<i>James M. Kaplan and Kayvaun Rowshankish</i>	
Chapter Three: Internet Openness and Fragmentation: Toward Measuring the Economic Effects	47
<i>Sarah Box</i>	
Chapter Four: Tracing the Economic Impact of Regulations on the Free Flow of Data and Data Localization	66
<i>Matthias Bauer, Martina F. Ferracane and Erik van der Marel</i>	
Section Three: Legal Jurisdiction and Interoperability	86
Chapter Five: Jurisdiction on the Internet: From Legal Arms Race to Transnational Cooperation	87
<i>Bertrand de La Chapelle and Paul Fehlinger</i>	
Chapter Six: Legal Interoperability as a Tool for Combatting Fragmentation	104
<i>Rolf H. Weber</i>	
Chapter Seven: A Primer on Globally Harmonizing Internet Jurisdiction and Regulations	114
<i>Michael Chertoff and Paul Rosenzweig</i>	
Section Four: Balancing Technical Openness and Fragmentation	119
Chapter Eight: Market-driven Challenges to Open Internet Standards	120
<i>Patrik Fältström</i>	
Chapter Nine: When Are Two Networks Better than One? Toward a Theory of Optimal Fragmentation.	131
<i>Christopher S. Yoo</i>	
Chapter Ten: A Framework for Understanding Internet Openness	141
<i>Jeremy West</i>	
About CIGI	151
About Chatham House	151
CIGI Masthead	151

ABOUT THE GLOBAL COMMISSION ON INTERNET GOVERNANCE

The Global Commission on Internet Governance was established in January 2014 to articulate and advance a strategic vision for the future of Internet governance. The two-year project conducted and supported independent research on Internet-related dimensions of global public policy, culminating in an official commission report — *One Internet*, published in June 2016 — that articulated concrete policy recommendations for the future of Internet governance. These recommendations address concerns about the stability, interoperability, security and resilience of the Internet ecosystem.

Launched by two independent global think tanks, the Centre for International Governance Innovation (CIGI) and Chatham House, the Global Commission on Internet Governance will help educate the wider public on the most effective ways to promote Internet access, while simultaneously championing the principles of freedom of expression and the free flow of ideas over the Internet.

The Global Commission on Internet Governance focuses on four key themes:

- enhancing governance legitimacy — including regulatory approaches and standards;
- stimulating economic innovation and growth — including critical Internet resources, infrastructure and competition policy;
- ensuring human rights online — including establishing the principle of technological neutrality for human rights, privacy and free expression; and
- avoiding systemic risk — including establishing norms regarding state conduct, cybercrime cooperation and non-proliferation, confidence-building measures and disarmament issues.

The goal of the Global Commission on Internet Governance is two-fold. First, it will encourage globally inclusive public discussions on the future of Internet governance. Second, through its comprehensive policy-oriented report, and the subsequent promotion of this final report, the Global Commission on Internet Governance will communicate its findings with senior stakeholders at key Internet governance events.

www.ourinternet.org

PREFACE

When I and my colleagues at the Centre for International Governance Innovation and Chatham House envisioned and launched the Global Commission on Internet Governance (GCIG) in 2014, we were determined to approach the work ahead strictly on the strength of evidence-based research. To make this possible, we commissioned nearly 50 research papers, which are now published online. We believe that this body of work represents the largest set of research materials on Internet governance to be currently available from any one source. We also believe that these materials, while they were essential to the GCIG's discussions over these past months, will also be invaluable to policy development for many years to come.

The GCIG was fortunate to have Professor Laura DeNardis as its director of research, who, along with Eric Jardine and Samantha Bradshaw at CIGI, collaborated on identifying and commissioning authors, arranging for peer review and guiding the papers through the publication process.

Questions about the governance of the Internet will be with us long into the future. The papers now collected in these volumes aim to be forward looking and to have continuing relevance as the issues they examine evolve. Nothing would please me and my fellow Commissioners more than to receive comments and suggestions from other experts in the field whose own research has been stimulated by these volumes.

The chapters you are about to read were written for non-expert netizens as well as for subject experts. To all of you, the message I bring from all of us involved with the GCIG is simple — be engaged. If we fail to engage with these key governance questions, we risk a future for our Internet that is disturbingly distant from the one we want.

Carl Bildt

Chair, GCIG

November 2016

**SECTION ONE:
A UNIVERSAL AND OPEN INTERNET**

INTRODUCTION: ONE INTERNET: AN EVIDENTIARY BASIS FOR POLICY MAKING ON INTERNET UNIVERSALITY AND FRAGMENTATION

Laura DeNardis

Copyright © 2016 by Laura DeNardis

ACRONYMS

APIs	application programming interfaces
AS	autonomous systems
CDNs	content delivery networks
DNS	Domain Name System
GCIG	Global Commission on Internet Governance
IDNs	internationalized domain names
IETF	Internet Engineering Task Force
IoT	Internet of Things
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ITU	International Telecommunication Union
IXPs	Internet exchange points
MLAT	Mutual Legal Assistance Treaty
NAT	network address translation
OECD	Organisation for Economic Co-operation and Development
TCP/IP	Transmission Control Protocol/Internet Protocol
W3C	World Wide Web Consortium

Two forces are in tension as the Internet evolves. One pushes toward interconnected common platforms; the other pulls toward fragmentation and proprietary alternatives.

– Kevin Werbach (2008)

The economic and social promise of bringing the next billion people online usually assumes the ongoing growth and availability of a universal Internet. But the Internet of the future has many possible trajectories. One twenty-first-century Internet policy debate concerns whether cyberspace will continue to expand into a single, universal network, or fragment into disjointed segments based on geographical borders or proprietary ecosystems. How this choice resolves in the contemporary context will have considerable implications for the future of global economic development, national security and counterterrorism, and for the nature of free expression and access to knowledge online.

The ability to interconnect a projected 50 billion objects — from health devices to industrial control systems — depends even more so on the pervasive interoperability and global reach afforded by the Internet, and the diffusion and integration of the network, far beyond mobile phones and laptops, deep into the everyday objects and infrastructures that support life’s day-to-day transactions. While the digital realm is still in its infancy, this *capacity* to connect ubiquitously to the Internet, regardless of location or access device, has become an implicit assumption of the twenty-first century.

Even in areas yet without Internet access, policy makers and entrepreneurs investing in information and communication technologies assume that building the necessary infrastructure is not only possible, but will empower citizens to participate in the global digital economy, access knowledge and engage in lawful communication with others, regardless of location or type of device. The more than 23,000 citizens polled in the 2014 CIGI-Ipsos Global Survey on Internet Security and Trust overwhelmingly view Internet access as a human right (see Figure 1), and vast majorities view the Internet as important for the future of free speech, political expression, access to knowledge, and to their economic well-being (CIGI-Ipsos 2014).


Eighty-three percent of users believe affordable access to the Internet should be a basic human right when asked: “How much do you agree or disagree with the following statement? ‘Affordable access to the Internet should be a basic human right.’”

In accord with these results, the United Nations Human Rights Council (2012) resolution on *The Promotion, Protection, and Enjoyment of Human Rights on the Internet* recognizes “the global and open nature of the Internet as a driving force in accelerating progress towards development in its various forms.”

That the growth and rapid technological development of the Internet, and access to it, now approaches a basic human right is a remarkable development given just how recently cyberspace and associated digital technologies have evolved. As Leslie Daigle, chief technology officer emerita of the Internet Society, has said, “A sign of success of the Internet is the degree to which we take it for granted” (Daigle 2014).

Not taking for granted the Internet’s interoperability and reach only requires recalling the computing environments that historically preceded it. Fragmentation was once patently the norm. Only a few decades ago, in 1981, IBM introduced its first personal computer. In the following decade, computer networks were disconnected isles of technology. Computers made by one company could be interconnected, but not with devices made by another. Digital networks were proprietary, based on closed technical specifications designed specifically *not* to connect

Figure 1: “How much do you agree or disagree with the following statement? ‘Affordable access to the Internet should be a basic human right.’”



Data Source: CIGI-Ipsos (2014).

with competitors’ products. Companies using one type of network, such as IBM’s Systems Network Architecture, could not communicate with a customer or business partner using a different environment, such as Digital Equipment Corporation’s DECnet or Apple’s AppleTalk network.

By design, there was no interoperability between systems. This architected lack of interconnectivity also characterized the popular, but proprietary, online consumer systems of the early 1990s, such as America Online, CompuServe and Prodigy, in which someone using one system could not communicate with someone using another. There was not yet interoperability — the ability to connect between devices, services and applications using standard protocols. The Internet, based on a family of protocols known as Transmission Control Protocol/Internet Protocol (TCP/IP), became the dominant open approach for enabling interconnectivity among diverse computing environments. The potential for universal reach and interoperability afforded by the Internet’s technical design was a significant departure from the proprietary and disjointed communication approaches of predecessor computer networks.

Some contemporary trends have raised concerns about movements back toward fragmentation. The revolutionary capacity for universal access and the aspirational expectations for the Internet’s accompanying economic and political benefits now stand in tension with geopolitical, technical and economic approaches poised to shift the Internet toward more of a segmented rather than universal system. Under the mantle of cyber sovereignty, governments have attempted to overlay geopolitical borders on the Internet, such as implementing efficient systems of content censorship and filtering, or enacting privacy-related laws mandating restrictions on where and how companies may store customer data. New business models, sometimes referred to as zero-rating services, designed to advance free access to the Internet in emerging markets, have raised questions about whether the next billion Internet users will

have access to the global Internet or to only a fraction of cyberspace available for free via walled gardens. There are also concerns about a resurgence of proprietary systems designed specifically not to interoperate with other systems, particularly in the context of new Internet of Things (IoT) products and services, but also as part of broad market trends away from general-purpose Internet access via browsers to mediation by platform-specific apps. These trends lead to the question of whether, over time, there will be a universal Internet or a fragmented Internet that varies based on country, region or proprietary ecosystem.

Conflicting values are always in tension in the realm of Internet architecture and governance — the broad ecosystem of administrative and design tasks necessary to keep the Internet operational — and the public-policy choices within this ecosystem. Tensions between network universality and enclosure indeed reflect conflicts regarding public-interest values in cyberspace, such as national security versus individual rights and freedom of expression versus privacy. They also reflect increasing incongruity between traditional Westphalian notions of sovereign nation states and a global technological system that crosses national borders and is overseen by a distributed, private-sector-led multi-stakeholder governance framework.

Objectives of national sovereignty and the global flow of information coexist tenuously. The coordination and technical design choices necessary to keep the Internet operational must constantly navigate diverging social values and interests. These alternatives are further complicated by the heterogeneous statutory, cultural and economic conditions that vary by region. To what extent should, or can, regional differences shape a distributed technical architecture that does not map neatly onto geographical borders?

Concern about Internet fragmentation emerged as a theme during the 2014 inception of the Global Commission on Internet Governance (GCIG). The commission viewed the

Internet governance debate about fragmentation not as the single issue so often portrayed in policy discourses, but as a constellation of questions crossing many layers of Internet infrastructure, involving many stakeholders, and with potential impacts that are not only technical, but economic and political. So began a process of commissioning scholarly work to examine various dimensions of Internet universality and Internet fragmentation, whether political, economic, infrastructural, legal or content-based. The objective of this research collection is to provide an analysis of the nature and implications of various forms of Internet fragmentation, with the ultimate purpose of improving the evidentiary basis of policy making in this area.

This introductory chapter helps to frame this research by, ironically, deconstructing (fragmenting) universal discussions about Internet fragmentation into a taxonomy of distinct topics that matches how the Internet works in practice and reflects the actual tangible policy choices at hand. Is there a universal Internet now? What are the various trends that could potentially move the Internet away from universality and toward fragmentation, and when is this desirable versus undesirable? What are the policy and design choices that can provide the capacity for a universal Internet but allow for institutional and individual freedom to not be completely interconnected? With these questions in mind, the following is divided into three sections:

- a consideration of the extent to which the contemporary Internet can be viewed as a universal network now;
- an exploration of the implications of emerging geopolitical and socio-economic initiatives associated with the potential for Internet fragmentation; and
- a baseline proposal for the technological characteristics and policy frameworks necessary for affording the Internet with a sustained capacity for ongoing global growth and openness.

THE STATE OF INTERNET UNIVERSALITY

Discussions about fragmentation frequently begin with the assumption that fragmentation is a new or emerging development that threatens the global reach and generativity of the Internet. At the level of infrastructure, the Internet is inherently a heterogeneous assemblage of thousands of different networks, primarily owned and operated by the private sector and able to interconnect only because they adhere to a common set of protocols specifying how to format and exchange information. Because of this interconnection and the capacity, generally, to move information from one point to another, regardless of geographical location, people speak of *the Internet*

and express concerns about whether it will fracture into *Internets*.

Examining the prospects and implications of Internet fragmentation first requires acknowledging that the contemporary Internet is not yet universal, geographically, materially or experientially. Divisions and barriers exist across the Internet ecosystem. Because of the complexity and heterogeneity of the global network, it can be useful to examine issues in layers, a conceptual framework that arose at least three decades ago around network protocols, such as the Open Systems Interconnection seven-layer protocol model (physical, data link, network, transport, session, presentation and application layers), or the more flexibly defined TCP/IP four-layer protocol suite (link, Internet, transport, application) (see, for example, Internet Engineering Task Force [IETF] 1989). This layered conceptual approach toward understanding protocols has given way to a norm of viewing the Internet as a layered system, even beyond protocols. In keeping with this tradition, which simply helps to conceptually organize the technological and administrative components of the Internet, this section will examine the state of Internet universality in four conceptual and overlapping categories:

- physical infrastructure (e.g., access, hardware, transmission systems);
- logical resources (e.g., IP addresses, protocols);
- the application and content layer (e.g., data and applications); and
- the legal layer (e.g., national policies and statutes, international treaties).

There is nothing fixed or natural about these categories, but they, or some variation of these categories, are frequently employed to discuss Internet architecture and policy, including discussions about fragmentation (Force Hill 2012; Drake, Cerf and Kleinwächter 2016). The layers also overlap with great complexity. For example, the legal (and policy) layer transcends the other three layers. Nevertheless, they are sufficient to help deconstruct the nuances of calling the Internet a universal network whose essential character may be threatened by fragmentation.

ASSESSING UNIVERSALITY AT THE PHYSICAL INFRASTRUCTURE LAYER

Viewed through the lens of physical infrastructure, the Internet is not yet a universal network. It must first be acknowledged that, by 2016, half of the world still does not have Internet access. According to International Telecommunication Union (ITU) indicators, 3.2 billion people had Internet access by 2015 (ITU 2015). Two billion of these users resided in developing countries, with many newer users accessing the network primarily from mobile

phones. Although half the world still does not have Internet access, the growth rate has been exorbitant. As recently as the year 2000, only 400 million people could access the Internet. This number has grown by 700 percent over a 15-year period.

Yet among the half of the world using the Internet, access speeds vary considerably. For example, broadband access speeds in countries such as South Korea, France, Iceland and Denmark are much faster, generally, than the speeds in countries throughout Africa and Latin America. There is also not an even distribution of Internet exchange points (IXPs) around the world, and almost half of countries do not have an IXP within their borders, although the IXP penetration rate is rising rapidly. IXPs are shared interconnection sites at which network operators make agreements to interconnect, thereby serving as essential nodes interconnecting the Internet's backbone. While access, interconnection penetration and access speeds vary, and while a digital divide persists, the trajectory historically has been toward greater access saturation, interconnection growth and broadband connection rates, all indicators of movement toward Internet universality.

ASSESSING UNIVERSALITY AT THE LOGICAL LAYER

Much of what keeps the Internet operational can be described as logical (meaning non-physical, virtual, or software-defined) resources. While the distinctions in practice are much more nuanced, general examples of the Internet's logical layer include: domain names; the global Internet address space of IP version 4 (IPv4) and IP version 6 (IPv6) binary numbers; the Domain Name System (DNS) that translates names into IP addresses; the thousands of protocols that standardize how information should be formatted, addressed, compressed, stored, encrypted, error-checked and transmitted over a network; and even architectural design principles, such as the "end-to-end" principle (Saltzer, Reed and Clark 1984). The end-to-end principle of locating intelligence at network end points has long been associated with the capacity for Internet universality. This groundbreaking technical design principle is often used to describe the logical structure of the Internet, but it does not always apply to the contemporary Internet because of the preponderance of "middle of the network" intelligence mechanisms, such as network address translation (NAT) and security firewalls.

There have historically been examples of fragmentation across all of these logical categories. For example, the Internet does not now have a completely universal address space because of the ongoing transition from one IP address standard to another. To exchange information over the Internet, each device uses a globally unique binary number, either permanently or temporarily assigned

for a session. The format of these IP addresses, under a long-standing protocol known as IPv4, assigns 32 bits to each binary address, a design choice that creates a global pool of 2^{32} , or roughly 4.3 billion Internet addresses. In the context of the internationalization and commercialization of the Internet, engineers anticipated that this would be an insufficient number to meet growth demands and designed a new standard, IPv6, to expand addresses to 128 bits long, providing an exponentially larger global address space of 2^{128} , or 340 undecillion addresses. For a variety of reasons related to political and economic incentives, as well as to technological complexities such as IPv6 being not natively backward-compatible with IPv4, IPv6 adoption has taken longer than anticipated (DeNardis 2009).

The Internet *had* a universal address space when the IPv4 address space was predominant, although even then some institutions used private address spaces on internal networks that connected to the global Internet through gateways. And it *would have* a universal address space if IPv6 adoption escalated to the point of deprecating IPv4. While the term "fragmentation" seems overstated, the Internet address space is not uniform in the contemporary context. This long-existing condition also produces, as Jonah Force Hill aptly describes, "serious interoperability problems within the crucial East/West Internet relationship" because the rate of IPv6 adoption across Asia, a place with far fewer IPv4 addresses than in the West, is so much higher than in the United States and Europe (Force Hill 2012). There is also sometimes fragmentation around the DNS when it is used to block local queries to certain websites, usually for content-blocking purposes such as censorship or enforcement of intellectual property rights.

ASSESSING UNIVERSALITY AT THE APPLICATION AND CONTENT LAYER

For those who do have access, the experience of Internet use varies considerably, often based on cultural and human-rights differences, such as what information is available in which language, level of digital literacy and what information is blocked or censored in a region. The spectrum of digital information available natively in English is much larger than the content available in other languages, so the experience of the Internet obviously varies based on language. Domain names, because they include content, have historically created language fragmentation. For most of the Internet's history, primarily because of its origin in the United States, domain names were only able to use the Latin alphabet, meaning that any languages using Arabic, Chinese, Cyrillic or other non-Latin characters were excluded from domain names. The standards community has developed the means to include non-Latin scripts via internationalized domain names (IDNs), but there are still barriers to the universal accommodation of these IDNs.

Fragmentation at the content level exists in part because of censorship. Information available online in China, in light of China's extensive system of filtering and blocking digital content, is quite distinct from the information available over the Internet in Sweden, for example. Fragmentation at the content level also arises from policies such as the "Right to be Forgotten" law in the European Union, which deletes content locally, or, in another example, geo-IP-restricted Netflix in Canada. The content available in one region is not necessarily the same as that content available in another region. With these differences in mind, the experience of the Internet at the content level is, of course, not universal.

There is also balkanization at the application level. Related to the diminishment of the end-to-end principle, most applications do not have the commensurable interoperability that existed with historically dominant Internet applications, such as email and the World Wide Web. With email, the expectation, and revolutionary innovation, was that anyone using an email client provided by one company could send emails to someone using a different email client. Similarly, someone could reach a website regardless of the browser or search engine used. Some contemporary Internet applications, ranging from Internet voice applications to social media to video games and messaging systems, do not have this interoperability, so are more fragmented. In the mobile environment, "apps" are tied directly to the platform provider and, often, the operating system and require platform mediation and curation. Some applications do not need to interoperate, or are designed not to interoperate for security reasons. For example, financial services applications often rely upon private networks or virtual private networks largely disconnected from the public Internet to achieve requisite performance metrics and security (Yoo 2016). But for general applications, taking the choice away from consumers to interoperate using common application types is a shift in norms. For example, there is no technical reason why making a voice call or sending a message over the Internet would require a proprietary system or gatekeeping function. It is a market technique. There is not necessarily interoperability among the apps used on different mobile platforms, either. Especially given the large number of users accessing the Internet via apps from mobile phones, this variation of fragmentation is significant.

Universal accessibility, however, has continuously improved for people with disabilities, such as those with sight or hearing impairments, largely because of the availability of Web accessibility standards established by the World Wide Web Consortium (W3C). Yet, despite gains, there are many opportunities for greater implementation of universal accessibility standards into applications.

ASSESSING UNIVERSALITY AT THE LEGAL LAYER

Although Internet governance is often viewed as one policy area, it is more accurately described as a broad ecosystem of tasks necessary to keep Internet technologies operational and the enactment of public policies around these technologies. The tasks are carried out by relatively new global institutions, such as the Internet Corporation for Assigned Names and Numbers and the IETF; the policies enacted by private Internet companies; international agreements; and national statutory and administrative frameworks. It is across this latter jurisdictional area of Internet governance that some of the greatest conflicts have historically arisen. The Internet is designed to be inherently cross-border, whereas national laws are bordered and vary significantly by jurisdiction in areas such as hate speech, privacy norms and approaches to intellectual property rights. Nation-state laws conflict with each other but especially stand in tension with the Internet's virtual, cross-border data flows and distributed character. Nations have jurisdictional oversight of the citizens and companies within their borders, but these borders do not comport well with the Internet's distributed and virtual nature.

Bertrand de La Chapelle and Paul Fehlinger (2016) warn about the implications of this disjuncture in their chapter *Jurisdiction on the Internet: From Legal Arms Race to Transnational Cooperation*. They argue that intergovernmental efforts fail to adequately address cross-border online challenges. Lacking is effective transnational cooperation, and national governments have undertaken legal and technical efforts to expand their jurisdiction in cyberspace. These efforts not only create international tensions, but also pose challenges to the stability of Internet infrastructure and human rights online. The authors recommend the creation of "issue-based governance networks" that facilitate transnational cooperation among actors based on shared principles which allow them to address issues such as requests for content removal.

In *Legal Interoperability as a Tool for Combatting Fragmentation*, Rolf H. Weber (2014) views legal interoperability as a means to prevent increasing Internet fragmentation and promote growth and expression online. Legal interoperability refers to the "process of making legal rules cooperate across jurisdictions" (ibid., 6). The extent to which legal mechanisms are balanced can be understood on a continuum, with complete assimilation and a fragmented legal landscape constituting the binary opposites. According to Weber, legal approaches need to be tailored to respective issues and contexts. A bottom-up approach is most effective in identifying legal solutions as it allows multiple stakeholders to come together to formulate solutions.

In the contemporary system, there is no harmonization of policy approaches across borders. In many cases, this is

preferable because legal harmonization toward repressive information policies would be problematic. In other cases, such as fighting cybercrime, greater cooperation would be desirable. The obvious challenge underpinning the question of legal harmonization is the question of jurisdiction — in other words, determining applicable laws in cross-border conflicts. Territoriality itself is difficult to assess because of complexities over whether jurisdiction is based on server location, user location, registrar location, or where a relevant intermediary is incorporated. While there are some legal treaties, such as the Council of Europe Convention on Cybercrime (also known as the Budapest Convention), there is still a great deal of diversity in legal approaches to the Internet, often shaped by political conceptions of what counts as freedom of expression and privacy and what is the appropriate role of the private sector. As such, cross-border requests have typically involved direct interactions between governments and private intermediaries, whether they entail user data requests, content blocking or another purpose. This approach presents challenges to information intermediaries, who have to navigate relevant and widely diverging laws in all the jurisdictions in which they operate, often under varying statutes regarding intermediary liability. Considering all of these factors, it cannot be said that there is a great deal of universality at the legal layer.

THE IMPLICATIONS OF EXOGENOUS TRENDS TOWARD FRAGMENTATION

While the preceding section indicates that various forms of fragmentation already exist throughout the Internet ecosystem, it also suggests that, especially at the infrastructure and logical layers, the Internet has continuously moved toward universality. Access rates continue to increase, IPv6 growth continues, new IXPs are built, IDNs are adopted. Policy and scholarly concerns about rising forms of Internet fragmentation have arisen from two exogenous trends around the Internet: market-driven fragmentation and geopolitically driven fragmentation. While it is also possible to create a separate discussion on purely technically driven fragmentation, the following section folds these technological issues into the discussions of economic and political contexts shaping Internet fragmentation, and then discusses the projected costs of fragmentation.

MARKET-DRIVEN FRAGMENTATION AND GEOPOLITICALLY DRIVEN FRAGMENTATION

Technological innovations such as the IoT and the rise in cloud-computing approaches create new spaces for the question of fragmentation versus universality. British computer scientist Dame Wendy Hall has said, “The Internet of Things is not yet an Internet.”¹ This is a prescient statement

because IoT implementations have not demonstrated, or aspired to, the same degree of interoperability and use of competition-enabling open standards as other areas of Internet applications. In *Market-driven Challenges to Open Internet Standards*, Internet engineer Patrik Fältström (2016) explains how market forces often oppose interoperability and competition in favour of locking users into proprietary services that are unable to interact with competitors’ services. This is particularly the case in emerging IoT markets. Fältström uses IP-based lighting-control systems as an example of both an IoT application and an emerging area in which manufacturers take non-interoperable, siloed approaches in which devices they manufacture speak to each other but not with devices made by other companies. These types of proprietary approaches that eschew interoperability and openness are the norm in consumer electronics, and, as Fältström explains, “each company imagines that its proprietary approach will become widely adopted as the ‘de facto’ standard, with respect to which it will have an obvious competitive advantage over other companies pursuing the same ‘maybe it will be me’ strategy” (ibid., 7). Another trend is the preponderance of cloud services in which users interact with the service via application programming interfaces (APIs) and are subject to the proprietary service’s terms and conditions rather than communicating based on standard protocols.

The question of market-driven fragmentation around technological disruption is part of a broader tension that has often arisen in the Internet space around private actors seeking market advantage through digital enclosure and proprietary approaches. In their white paper on Internet fragmentation produced for the World Economic Forum’s Future of the Internet Initiative, William Drake, Vinton Cerf and Wolfgang Kleinwächter (2016) provide an extensive taxonomy of the types of commercially driven fragmentation that occur, including peering and interconnection, certain types of net-neutrality violations, walled gardens and geo-blocking of content.

Rising geopolitical challenges around the Internet similarly raise concerns about the prospects for a universal Internet. Jurisdictional conflicts that have always accompanied Internet globalization are complicated by emerging economic, political and technical factors. The economic stakes of digital commerce are high, political contention over content control is rising, and technological structures — such as cloud computing and content distribution networks — are increasingly distributed. More than ever, technologies do not reside neatly within borders, and therefore jurisdictions. Where data is stored (often in more than one place via replication and caching), where a domain name is registered, where employees reside and where a company is incorporated no longer have natural relationships.

In this context of blurred lines between technological and national borders, some governmental policies are seeking

1 Personal communication to author.

to reassert geographical sovereignty in cyberspace, often in specific policy areas. Data localization laws are a prime example. These laws place constraints on how and where private companies store customer data, such as requiring customer data to be stored on servers within a nation's borders or placing various restrictions on the nature of and extent to which customer information is shared across borders (Chander and Le 2015). The impetus for some of these policies concerns customer privacy in the context of foreign surveillance. Accordingly, some arose in the contentious aftermath of disclosures about the expansiveness of the surveillance program of America's National Security Agency. In other cases, the motivation is to create market advantages for indigenous rather than foreign companies.

Data localization laws raise many questions about potential effects on engineering efficiency, the cost of doing business, the ability to innovate and human rights. Concentrating data in a fixed location can actually facilitate efficient surveillance, either from the host country or via foreign surveillance. From an engineering perspective, factors that affect how information is stored and transmitted include the goals of reducing latency, providing redundancy and replication to distribute data closer to its destination, and other basic traffic-engineering and traffic-optimization goals that can conflict with data localization requirements.

Politically driven infrastructure prescriptions also heighten concerns about legal fragmentation. In *A Primer on Globally Harmonizing Internet Jurisdiction and Regulations*, Michael Chertoff and Paul Rosenzweig (2015, 1) warn about the potential legal fracturing of the Internet due to geopolitical trends such as data-localization policies: "We stand on the cusp of a defining moment for the Internet. Existing trends, left unaddressed, might very well lead to the fracturing of the World Wide Web."

Their chapter extends the question of which nations' laws jurisdictionally apply in different contexts. In other words, who has power over what? As an alternative to the jurisdictional concerns raised in data localization laws, Chertoff and Rosenzweig propose and evaluate a choice-of-law rule based on four models for clarifying jurisdiction: citizenship of data creator, citizenship of data subject, location of "harm" that has taken place, or citizenship of data custodian. They also provide recommendations about streamlining the Mutual Legal Assistance Treaty (MLAT) structure, which could help minimize incentives for unilateral approaches such as data localization rules.

THE ECONOMIC EFFECTS OF OPENNESS AND FRAGMENTATION

Discussions about the effects of infrastructure prescriptions such as data localization laws often centre on large content intermediaries like Google. What is often overlooked is that these laws also have significant effects on other

economic sectors. From financial services to retail, every sector of the economy relies upon digital technologies to store and transmit information about customers or engage in routine business practices such as billing or the delivery of services. Similar to the tech sector, many of these companies in other industries have customers, stores and offices throughout the world, and are not concentrated in any particular country.

James M. Kaplan and Kayvaun Rowshankish (2015) of McKinsey & Company address the economic implications of data localization laws on the financial services sector in their chapter *Addressing the Impact of Data Location Regulation in Financial Services*. Their survey of chief executives in the financial-services sector suggests that data localization laws place significant burdens on private industry, including the complexity costs of navigating and interpreting different regulations across jurisdictions, and of either making technological modifications to comply with new regulations or pulling out of certain markets entirely. For example, to comply with some laws, financial-services companies must locate human resources and technical infrastructure in places where they otherwise would not have a physical presence. As they explain, "Data location regulations make some countries economically unattractive, causing institutions to exit, and limiting their global footprint" (Kaplan and Rowshankish 2015, 3).

The Organisation for Economic Co-operation and Development (OECD) has been doing work to measure global data flows and quantitatively assess the effects of Internet openness. In her chapter *Internet Openness and Fragmentation: Toward Measuring the Economic Effects*, OECD senior policy analyst Sarah Box (2016) presents some of the initial results and, in particular, OECD efforts to aggregate and analyze cross-border data flows among the world's countries using corporate data from Google searches and YouTube views. A universally accessible Internet that enables free flows of information across borders is widely understood to have positive effects on trade, whether by improving supply-chain efficiency, expanding customer and market reach, or bettering payment and delivery systems. The knowledge shared freely across borders also stimulates innovation and entrepreneurship. Box's chapter addresses the difficulty of establishing empirical evidence of these connections, describes some of the existing studies quantifying the effects of Internet openness, and presents some of the OECD's initial findings, including a "uniform trend of users increasingly accessing content outside their countries," as well as establishing that data flows, while not predictable, often have international dimensions (*ibid.*, 6).

Laws that limit the free flow of information across borders have detrimental effects on the wider economy beyond implications to industry. Researchers Matthias Bauer, Martina F. Ferracane and Erik van der Marel (2016) quantitatively present the broad costs of data localization laws in their study *Tracing the Economic Impact of Regulations*

on the *Free Flow of Data and Data Localization*. They developed an index that serves as a proxy for data regulation across various OECD and emerging economies, and then assess the impact of regulations on downstream sectors that make use of data. Their study examines specific laws in 60 jurisdictions, and quantitatively models how data localization laws would engender losses to GDP, decreases in domestic investments and welfare losses to citizens. They conclude, “Accordingly, tight regulations on the free flow of data tend to cause an economy’s production structure to shift (back) towards less innovative and relatively volatile sectors such as agriculture, raw materials and natural resources” (ibid., 18).

Another dimension of analysis is that bordered Internet policies rarely correspond to how Internet infrastructure works in practice. Although physical infrastructure such as fibre-optic cable, switching centres, routers and radiofrequency antennas reside within physical borders, neither the logical architecture nor the realities of how information flows over the Internet comport neatly with national borders. This is especially the case in interconnection issues. Routers make decisions about how to forward packets based on issues of network efficiency and resource reachability rather than on where the next hop physically resides. The actual “bordered” areas of the Internet are autonomous systems (AS). The Internet is described as a network of networks but it is more technically accurate to describe it as an interconnected network of virtual AS. Autonomous systems are routing domains, which manage a set of IP addresses either residing in the domain or accessible through that domain to an entity that pays a transit fee to connect to the global Internet through that system. Most understand that handoffs between network operators also require physical interconnections, such as those that occur at shared IXPs. But even these interconnection points do not correspond to a geopolitically bordered view of the Internet, because an exchange of information originating and terminating between two telecommunication companies within a single country can potentially be routed through an IXP located in another country, before being routed back to the originating region.

How company business models, across all sectors of the economy, also use the Internet does not correspond to national borders. Companies can register a domain name in one country; locate servers in another; establish customer service centres in yet another country; and hire content delivery networks (CDNs) or cloud-computing providers to replicate, store or cache information all over the world. Geopolitically driven policies that seek to place borders around dimensions of Internet data flows should also consider the intractability of aligning these policies with the material and virtual reality of how the Internet actually works.

A TECHNICAL DESIGN AND POLICY VISION FOR A UNIVERSAL INTERNET

Internet governance is not static any more than the Internet’s technical architecture is static. Contemporary policy choices will affect not only a spectrum of public-interest issues but also the stability and character of the Internet itself, in the same way that architecture reciprocally shapes policy choices. Although various forms of fragmentation already permeate the Internet ecosystem, the generative and open qualities of the network have nevertheless enabled its rapid geographical expansion, and have also created conditions that generally promote an open playing field for entrepreneurs to introduce new systems and applications that could be assured to interoperate with other systems globally. There has been diversity in the types of devices, services and applications enabled largely by conformance to open technical protocols that allow these diverse environments to exchange information with each other.

Given that technological change has been constant in the Internet environment, what fundamental principles or other design characteristics have enabled this growth and innovation? Internet engineer Leslie Daigle (2015), in her chapter *On the Nature of the Internet*, acknowledges the constant and rapid transformations in the Internet’s underlying technical architecture and suggests that it may be preferable to define the Internet based on its core underlying principles, or “Internet invariants,” as the Internet Society (2012) has described these characteristics. These principles include: global reach/integrity; general purpose; supporting innovation without requiring permission; accessibility; interoperability and mutual agreement; collaboration; reusable (technical) building blocks; and no permanent favourites (Daigle 2015).

All of these principles speak in some way to the Internet’s inherent potential for universality. For example, the principle of *global reach* is designed to allow any two devices connected to the Internet to connect with each other, regardless of location or network. The diversity principle of *general purpose* expands this goal to allowing for any application or service to run over the Internet. The principle of *permissionless innovation*, the ability for anyone to set up a new service without requiring anyone else’s permission, is linked closely to the universality and openness of the Internet because it creates the capacity and potentiality of innovation to arise from anywhere in the world, and without having to pass through gatekeeping constraints. A closely related principle is *no permanent favourites*. Because the Internet’s underlying technical infrastructure enables anyone to connect and introduce new innovation, new entrants are always possible and, in a continuous cycle of disruption and innovation, the entrepreneurs of today are potentially the dominant business people of tomorrow. Perhaps most salient to the potential of a universal Internet is the principle of *interoperability and mutual agreement*.

What has operationalized many of the principles leading to the capacity for Internet universality are the open technical standards that are developed collaboratively in standards-setting institutions such as the IETF and the W3C, as noted, and made publicly available so that others can develop products with the assurance of compatibility with heterogeneous services, devices and applications on the Internet. Internet standards serve as the blueprints developers can use to ensure that their products are interoperable with other products in the marketplace. These standards serve a primary technical purpose, but they also carry political implications and economic externalities. Politically, these institutions sometimes make public-interest decisions, such as on the extent of user privacy or accessibility for the disabled. Economically, technical standards, and the extent to which they have embedded intellectual-property restrictions, are closely linked to innovation because they provide a platform upon which innovation and competition can occur (DeNardis 2011).

Open standards are therefore linked to the question of Internet universality versus fragmentation in three ways. If technical standards sometimes establish public policy, procedural norms of participatory openness, as well as open publication of the standard, are necessary to establish policy-making legitimacy; technically, they provide the interoperability among applications, networks, and services that is necessary for the possibility of global accessibility and reach; and economically, open standards are the primary enabler of market competition and the operationalization of the innovation principle of no permanent favourites.

At the same time, network fragmentation does not always produce detrimental effects. Many of the core technologies necessary for cyber security and basic business operations, such as firewalls and virtual private networks, are designed precisely to “fragment” the Internet. A network with sensitive health records or financial data should not be universally accessible or interoperable. In his chapter *When Are Two Networks Better Than One? Toward a Theory of Optimal Fragmentation*, Christopher S. Yoo (2016) references Metcalfe’s law concerning the value of connectivity based on the network-effect insight that, as a network grows, accretion in the number of connections exceeds the growth in the number of nodes. After a point, there can be diminishing marginal returns with additional resources on a network. Yoo also notes that concern about fragmentation must take into account not just optimization of the network as a whole, but also incentives for individual actors.

Lack of interconnection, interoperability and universality are sometimes beneficial, and are indeed carefully designed into systems for the purpose of securing private communication systems or carefully controlling access to and from the global public Internet. But this is an example of a design choice applied to a private network that private entities should be allowed to make, in the same way they should be allowed to choose to connect their private networks to

the global public network. Choosing to limit connectivity in certain ways does not foreclose the possibility of connecting in the future or under different circumstances. The potential choice of openness is indeed part of openness.

Many contemporary forces are in tension with traditions of openness: market-driven approaches that seek enclosure and proprietary advantage; geopolitically driven policies that seek to place borders on the Internet; lack of adoption of technologies that address digital resource constraints; and various types of content fragmentation, ranging from censorship to infrastructure-based, intellectual-property-rights enforcement. It is also clear that forces seeking to move the Internet toward greater fragmentation come from both government and the private sector, all complicated by technological disruptions. Furthermore, user choices, to some extent, are also selecting approaches that are arguably more fragmented, such as widespread adoption of proprietary and non-interoperable social-media applications and messaging systems. A great question is whether these tensions will have long-term detrimental effects on the character of the open Internet.

Of course, it has become a mantra to express that the Internet should remain “free and open.” But defining “free” and “open” is difficult in practice. Open-source-software communities often make the distinction between “free beer” and “free speech.” So too, openness in the context of Internet governance is contextual and can refer to technical openness (open standards), civil-liberties openness (freedom of expression and association), and openness of digital markets (permissionless innovation and a level playing field for competition). When the term “Internet openness” is used, it can take on any or all of these meanings.

In *A Framework for Understanding Internet Openness*, OECD senior policy analyst Jeremy West (2016) seeks to answer the enigmatic question of what Internet openness is. West posits that there is “no such thing as *the* open Internet,” but rather, “Internet openness, which exists in various degrees along several dimensions” (2016, 1) and that “the essence of Internet openness is the global free flow of data across the network” (ibid., 8). The OECD’s ongoing work on Internet openness has helped advance an understanding that accounts for network and social heterogeneity while defining openness at three levels: technical, economic and social. Technical openness refers primarily to features of interoperability and universality, such as a universal address space, open protocols and inclusive technology governance. Economic openness refers to features such as infrastructure access at a competitive cost, the capacity for cross-border digital exchange, and regulatory transparency and certainty. Social openness invokes a collection of human rights online, such as the right to privacy, the right to education, and rights of freedom of expression and association.

Table 1: Baseline Characteristics of Internet Universality

Layer	Internet Governance Characteristic
Physical Infrastructure	<ul style="list-style-type: none"> • Investments in broadband access penetration • Policies that promote the development of IXPs and other interconnection and transmission systems in emerging markets • Human capacity building
Logical Resources	<ul style="list-style-type: none"> • A universal IP address space • A universally consistent and stable DNS • Adoption of IPv6 • Open technical standards that are open in participation and implementation, and engender multiple competing products that are interoperable • Human capacity building in standards setting and critical logical resources
Application and Content Layer	<ul style="list-style-type: none"> • Promotion of global access to knowledge rather than censorship of lawful content • Universal support of IDNs • Applications that adopt standards of accessibility for the disabled • Promotion of digital literacy • Promotion of interoperability norms in emerging contexts such as IoT
Legal Layer	<ul style="list-style-type: none"> • Rejection of government policies that restrict the flow of data across borders and have detrimental effects on trade, economic growth and freedom of expression • Agreements among governments to not tamper with the core infrastructure of the Internet, such as the DNS and systems of routing and interconnection • Promotion of the private-sector-led multi-stakeholder governance system

Source: Author.

This collection, taken as a whole, advances research and informs policy making in several ways. It suggests that, while the Internet has not yet achieved universality, its aspirational capacity for global reach and interoperability is being challenged by a number of exogenous pressures, both market-driven and geopolitical. Systems of Internet infrastructure and governance are increasingly recognized as critical points of control for achieving market advantage or carrying out geopolitical or global economic objectives. Many efforts to gain political and economic advantage bring the network toward fragmentation and away from universality, and this movement is not without costs to national economies, human rights, and the stability and security of the Internet. Preserving one Internet requires policies (see Table 1) that: incentivize infrastructure advancements such as the adoption of IPv6, growth in broadband access, and the global distribution of IXPs and undersea cables; promote trust by providing strong cyber security and a universal framework of basic human rights online; promote conditions for open innovation models geared toward permissionless innovation and access to knowledge rather than proprietary advantage and information enclosure; and preserve the inclusive and participatory multi-stakeholder model of Internet governance over emerging efforts geared toward cyber sovereignty, multilateralism and state control. As Internet technological disruption rapidly evolves toward the IoT and other emerging cyber systems pervading every corner of social and economic life, the enclosure or openness of

these new market innovations will help determine whether the digital sphere is constituted by non-interoperable fragments or a universal Internet.

WORKS CITED

- Bauer, Matthias, Martina Ferracane and Erik van der Marel. 2016. *Tracing the Economic Impact of Regulations on the Free Flow of Data and Data Localization*. GCIG Papers Series No. 30. Waterloo, ON: CIGI. www.cigionline.org/publications/tracing-economic-impact-of-regulations-free-flow-of-data-and-data-localization.
- Box, Sarah. 2016. *Internet Openness and Fragmentation: Toward Measuring the Economic Effects*. GCIG Papers Series No. 36. Waterloo, ON: CIGI. www.cigionline.org/publications/internet-openness-and-fragmentation-toward-measuring-economic-effects.
- Chander, Anupam and Uyen Le. 2015. "Data Nationalism." *Emory Law Journal* 64 (3): 677-739.
- Chertoff, Michael and Paul Rosenzweig. 2015. *A Primer on Globally Harmonizing Internet Jurisdiction and Regulations*. GCIG Paper Series No. 10. Waterloo, ON: CIGI. ourinternet.org/publication/a-primer-on-globally-harmonizing-internet-jurisdiction-and-regulations.
- CIGI-Ipsos. 2014. CIGI-Ipsos Survey on Global Security and Trust. www.cigionline.org/internet-survey.

- Daigle, Leslie. 2014. "Permissionless Innovation — Openness, Not Anarchy." *Internet Society Tech Matters* (blog), April 22. www.internetsociety.org/blog/tech-matters/2014/04/permissionless-innovation-openness-not-anarchy.
- . 2015. *On the Nature of the Internet*. GCIG Paper Series No. 7. Waterloo, ON: CIGI. www.ourinternet.org/publication/on-the-nature-of-the-internet.
- de La Chapelle, Bertrand and Paul Fehlinger. 2016. *Jurisdiction on the Internet: From Legal Arms Race to Transnational Cooperation*. GCIG Paper Series No. 28. Waterloo, ON: CIGI. www.ourinternet.org/publication/jurisdiction-on-the-internet.
- DeNardis, Laura. 2009. *Protocol Politics: The Globalization of Internet Governance*. Cambridge, MA: MIT Press.
- , ed. 2011. *Opening Standards: The Global Politics of Interoperability*. Cambridge, MA: MIT Press.
- Drake, William J., Vinton G. Cerf and Wolfgang Kleinwächter. 2016. "Internet Fragmentation: An Overview." World Economic Forum Future of the Internet Initiative White Paper, January. www3.weforum.org/docs/WEF_FII_Internet_Fragmentation_An_Overview_2016.pdf.
- Fältström, Patrik. 2016. *Market-driven Challenges to Open Internet Standards*. GCIG Papers Series No. 33. Waterloo, ON: CIGI. www.cigionline.org/publications/market-driven-challenges-open-internet-standards.
- Force Hill, Jonah. 2012. "Internet Fragmentation: Highlighting the Major Technical, Governance and Diplomatic Challenges for U.S. Policy Makers." Belfer Center for Science and International Affairs, John F. Kennedy School of Government, Harvard University. http://belfercenter.hks.harvard.edu/files/internet_fragmentation_jonah_hill.pdf.
- IETF. 1989. "Request for Comments 1122: Requirements for Internet Hosts — Communication Layers." Edited by Robert Braden, October. www.tools.ietf.org/html/rfc1122.
- Internet Society. 2012. "Internet Invariants: What Really Matters." www.internetsociety.org/internet-invariants-what-really-matters.
- ITU. 2015. "ICT Facts and Figures — The World in 2015." www.itu.int/en/ITU-D/Statistics/Pages/facts/default.aspx.
- Kaplan, James M. and Kayvaun Rowshankish. 2015. *Addressing the Impact of Data Location Regulation in Financial Services*. GCIG Paper Series No. 14. Waterloo, ON: CIGI. www.ourinternet.org/publication/addressing-the-impact-of-data-location-regulation-in-financial-services.
- Saltzer, Jerome, David Reed and David Clark. 1984. "End-to-End Arguments in System Design." *ACM Transactions on Computer Systems* 2 (4): 277–88.
- United Nations Human Rights Council. 2012. *Resolution on the Promotion, Protection and Enjoyment of Human Rights on the Internet*. www.ohchr.org/EN/HRBodies/HRC/RegularSessions/Session20/Pages/ResDecStat.aspx.
- Weber, Rolf H. 2014. *Legal Interoperability as a Tool for Combatting Fragmentation*. GCIG Papers Series No. 4. Waterloo, ON: CIGI. www.ourinternet.org/publication/legal-interoperability-as-a-tool-for-combatting-fragmentation.
- Werbach, Kevin. 2008. "The Centripetal Network: How the Internet Holds Itself Together, and the Forces Tearing it Apart." *University of California Davis Law Review* 42: 343–412.
- West, Jeremy. 2016. *A Framework for Understanding Internet Openness*. GCIG Papers Series No. 35. Waterloo, ON: CIGI. www.cigionline.org/publications/framework-understanding-internet-openness.
- Yoo, Christopher S. 2016. *When Are Two Networks Better than One? Toward a Theory of Optimal Fragmentation*. GCIG Papers Series No. 37. Waterloo, ON: CIGI. www.cigionline.org/publications/when-are-two-networks-better-one-toward-theory-of-optimal-fragmentation.

ABOUT THE AUTHOR

Laura DeNardis, CIGI senior fellow, is a scholar of Internet architecture and governance and professor in the School of Communication at American University in Washington, DC. The author of *The Global War for Internet Governance* (Yale University Press, 2014) and several other books, her expertise has been featured in numerous publications. She serves as the director of research for the Global Commission on Internet Governance and is an affiliated fellow of the Yale Law School Information Society Project, where she previously served as executive director. Laura holds an A.B. in engineering science from Dartmouth College, a Master of Engineering from Cornell University, a Ph.D. in science and technology studies from Virginia Tech, and was awarded a postdoctoral fellowship from Yale Law School.

CHAPTER ONE: ON THE NATURE OF THE INTERNET

Leslie Daigle

Copyright © 2015 by Leslie Daigle

ACRONYMS

ASN	Autonomous System Number
ASs	autonomous systems
BGP	Border Gateway Protocol
DNS	Domain Name Service
ETNO	European Telecommunications Network Operator
HTTP	HyperText Transmission Protocol
HTML	HyperText Markup Language
IANA	Internet Assigned Number Authority
ICE	Immigration and Customs Enforcement (US)
IETF	Internet Engineering Task Force
IMAP	Internet Mail Access Protocol
IP	Internet Protocol
IPv4	IP version 4
IPv6	IP version 6
ISP	Internet Service Provider
IXPs	Internet eXchange Points
NATs	Network Address Translators
NTP	Network Time Protocol
PIPA	Protect IP Act
RFC	Request for Comments
RIRs	Regional Internet Registries
SMTP	Standard Message Transmission Protocol
SOPA	Stop Online Piracy Act
TLD	top-level domain
WWW	World Wide Web

INTRODUCTION

A firm grasp of the nature of the Internet is required to help chart its future through the integration of policy and technology world views. There are many complexities — in technology and in the policy and use of the Internet — that can be difficult to characterize accurately as either key issues or passing distractions. This chapter describes the nature of the Internet with a view to furthering an understanding of the relationship between policy and technology, and how policy can help or hinder the Internet.

The Internet is no stranger to massive change. It is vastly different today from how it was at its inception — that the Internet has evolved over the course of 40-plus years is a testament to its flexibility in the face of major change. Over the years, however, there have been various predictions of technical causes of impending doom for the

network.¹ The reasons for concern were real, but crisis was averted through some explicit or implicit collective action. Additionally, some of the disastrous outcomes have been avoided by incremental degradation of the overall system known as the Internet.²

As the Internet and the services it supports continue to become an integral part of personal, commercial and political daily lives, there are increasing non-technical pressures on the Internet. There is perceived need for change in the Internet, often met by resistance from key stakeholders. Yet the Internet must be able to withstand some changes without losing its core nature — indeed, change is how the Internet has grown.

The Internet’s technical community, responsible for the development, deployment and operation of the Internet, and the world’s policy makers, responsible for the care of their citizens on- and offline, have increasingly found themselves in heated discussion over how to address policy issues without “breaking” the Internet. In the worst case, policies imposed on network operators, content providers and users of the Internet do not work (fail to address the issue for which the policy was created) and stifle the Internet’s growth and evolution. Sometimes, the policy measures succeed but the Internet’s growth is stifled — leaving the technical community wishing that different approaches could have been brought to bear. Or, the policy issue is not addressed, leaving policy makers and regulators unsatisfied and with ongoing concerns. None of these outcomes is particularly desirable. To make steps toward the ideal outcome (policy issue addressed and Internet’s growth unimpeded), a broader understanding of the nature of the Internet is needed, without requiring policy makers to be ready to argue technical points or vice versa.

How can one distinguish between helpful and healthy adjustments to the Internet and actions that will undermine the nature of the Internet? How can one engage in meaningful dialogue across stakeholders, including those more versed in how the Internet works and those who understand the needs of the world’s communities?

1 For example, in 1995, Ethernet inventor and industry leader Bob Metcalfe famously said, “I predict the Internet will soon go spectacularly supernova and in 1996 catastrophically collapse.” It did not, and he literally ate his own words in the form of a blenderized copy of his printed prediction paper, at the Sixth International World Wide Web Conference in 1997 (Goble 2012).

2 “Network Address Translation” was introduced to allow several computers to share a single external Internet Protocol (IP) address, in the face of IP version 4 (IPv4) addresses becoming scarce. However, this means that those computers are not directly reachable on the Internet, since the address is handled by a gateway box that serves several computers at once.

Key to answering those questions is understanding the nature of the Internet in terms that are not strictly technical. This chapter will:

- outline the technical nature of the Internet;
- articulate the unchanging properties of the Internet (the “invariants”); and
- leverage both of those frameworks to examine current challenges facing the Internet.

The concerns for change are not strictly hypothetical. The Internet is currently facing several situational challenges. There are proposed (and some implemented) policies in the world that are meant to address very real concerns, but that negatively impact the Internet’s operation, growth and value as a platform for continued innovation. This chapter will review, through the lens of the Internet’s invariant properties, various challenges the Internet is currently facing.

THE TECHNICAL NATURE OF THE INTERNET

This section provides a general overview of Internet technology as a necessary background for understanding key points in the rest of the chapter. It is intentionally high level, aiming to underscore key aspects of technology rather than attempt a complete exposition. Readers who are familiar with Internet technology may prefer to skim the section for key points of focus.

NETWORKS

In simplest terms, a network is something that connects different participants. In the context of the Internet, these participants have traditionally been called hosts. Initially, hosts were typically large-scale computers, on the scale of mainframes and then minicomputers. Gradually, as computing power increased, computing devices got smaller and more specialized. These days, just about anything can be a “participant” in an Internet network — everything from large computers to desktops to notebooks to mobile phones and car components.

“Connecting” participants means different things in disparate networks. For telecommunications networks, connection is providing a means to communicate between participants. Where telecommunications networks differ is in terms of their approaches to identifying participants, managing passage of information between those participants and the types of communications enabled within the network. For example, traditional telephony networks in the twentieth century used telephone numbers to identify endpoints, country codes and within-country area codes to find the phone being called, and established connections between participating telephones in order to

enable voice communication over the established channel. The rest of this section provides more detail on how the Internet generation of networks identifies participants and other details. At its inception, the Internet distinguished itself from traditional telecommunications networks by taking the approach of “connection-less” management of information passage. Unlike the traditional telephone network, information passage is achieved by carving up the information and putting “chunks” of data into “packets.” These packets contain all the necessary information to specify the intended destination and no information about required paths. Packets are sent independently through the network, over whatever channels work best at that instant in time.

PROTOCOLS

Standards are required in order to connect participant hosts from every manufacturer, all over the world, in all networks. These standards define everything from the expected voltages and electrical requirements of physical network hardware to the higher level of information exchange needed to carry out human communications. When it comes to standardizing the communication between Internet hosts — from the basics of passing packets of data to the more involved communications between end-users of the network — the standards define *protocols*. Protocols are the rules of the road, the lingua franca of Internet communications. The IP defines the layout of the individual packets of data mentioned above. This standard provides the definition that allows receiving hosts to “read” the packets (determine where the packet came from, where the bits of data “payload” are and so on), and it defines how sending hosts should form valid packets for transmission on the Internet. Within the IP packets, the data payload is not just a jumble of bits. Rather, it is structured according to the standard defined for some higher-level (closer to the end-user) protocol — for example, it might be part of a communication with an email server and governed by the protocol for interacting with that type of server.

INTERNET ADDRESSES

While the protocols define the rules of the road for communications on the Internet, the hosts are identified by addresses. Every host (machine, phone or component of a car) that is on the Internet is assigned a unique address when it connects to the Internet — a unique IP address. One host connecting to another on the Internet uses the IP standard to create packets, including its own IP address and the address of the destination host within each packet. As such, IP addresses are critical to maintaining a global, growing Internet. The version of the IP standard that is most commonly in use today is IPv4. Twenty years ago, it was apparent that the growth of the Internet beyond the purposes of academic research meant that the number of

unique addresses available in IPv4 — roughly four billion — would not be adequate to provide a unique address to every host on the Internet. After all, there are more people on the planet than there are IPv4 addresses. IP version 6 (IPv6) was standardized, with vastly more addresses available, and it is now being increasingly deployed to ensure global connectivity.

MOVING PACKETS: ROUTING

Once the source and destination addresses are known, there is still work to be done to get a packet from the origin host to its destination: routing. There is some merit in considering an analogy for routing: “turn-by-turn navigation” in modern GPS devices. Five cars (packets) may set out from one home (origin host) and travel different, but possibly overlapping, paths (routes) to a restaurant (destination host). Depending on the time of day, traffic on the road or other considerations, different choices in routing may be made. The process is a little different if you are going to a restaurant in a different town. You might first drive to the other town (on your generally preferred highway, or on the scenic route through a picturesque landscape and small towns) before turning on the GPS to find the exact location of the restaurant.

The useful points of analogy include the fact that choices are made based on current conditions and preferences. It is not that there are exactly five paths from the house to the restaurant, but rather that there are many possibilities and choices made for each segment, resulting in variations in path taken. Also, the notion of first working out how to get to a general vicinity and then using a more refined means of location also applies.

The analogy does fall apart if you press into how routes are determined in GPS navigation versus internetworking, so take the analogy for what it is.

As an internetwork, routing of Internet traffic happens to get a packet from one network to another, which may or may not be directly connected. Routes are advertised within the routing system — one network will share its path and connectivity to certain other networks. Based on these advertisements, packets will be forwarded through and between networks to reach a final destination network.

NETWORK BOUNDARIES OR EDGES

There are boundaries on networks: generally, a network is under one entity’s control (Internet Service Provider [ISP], enterprise, government or other form of public or private operator). But one entity may operate multiple networks, or at least provide multiple network faces to the rest of the world. Each such face, or routing unit, is an autonomous system and is identified in the routing system by an Autonomous System Number (ASN). These ASNs, the allocation of which is managed by the Regional Internet

Registries (RIRs), are the basis of the identification of paths through the Internet.

The important thing to note about these ASs is that they have boundaries and topology in a network sense, not a geographic sense. While they may be contained in a warehouse of servers, or spread across vast swathes of physical geography, the geography they cover may be unique to that network or there might be multiple networks crossing the same space: each AS is its own world.

CONNECTING NETWORKS

In order to have a global network then, these autonomous networks need to be hooked up — internetworked. This is done by creating gateways between networks — where a network router is set up to take traffic that is destined for hosts outside the network and pass it to a neighbouring network for onward transmission, or accept incoming traffic from a neighbouring network and route it internally. In order to manage these connections between networks, the Border Gateway Protocol (BGP) standard is used (Rekhter, Li and Hares 2006).

BGP is *how* routers communicate to connect networks. Agreements between network operators determine which networks are connected and the policies under which network traffic will be carried. Operators may choose to connect as “peers” (peering). In the case of large networks, where there is symmetry in the amount of traffic that each would send to or through the other network, this might be done on a cost-free basis. Otherwise, a smaller network may “buy transit” from a larger network, paying to connect to the larger network in order to get access, or better access, to relevant parts of the Internet. A more recent popular alternative is for networks to connect to so-called Internet eXchange Points (IXPs), where they can exchange traffic directly with other networks at the IXP and not have to pay for upstream transit of the traffic. This makes it possible to “keep local traffic local.”

APPLICATIONS AND SERVICES INFRASTRUCTURE

Of course, the Internet requires more than just connections between networks in order to support the key uses the world has come to know and depend on. Internet applications are built as software to implement application protocol standards. Electronic mail, or email, is transmitted through one standard protocol, Standard Message Transmission Protocol (SMTP) (Klensin 2008), and can be retrieved from servers using a different standard protocol, such as the Internet Mail Access Protocol (IMAP) (Crispin 2003). As originally conceived, every host on the Internet was expected to run a mail server program that could send and receive mail messages. In practice, this led to a lot of spam messages being sent via “open relay” mail servers, and it became more common for household customers

of ISPs to send mail through their ISP's mail servers. The World Wide Web (WWW) is another Internet application — clients connect to WWW servers using the HyperText Transmission Protocol (HTTP) (Fielding and Reschke 2014).

None of the above would be especially useful without the Domain Name Service (DNS) standard protocol (Mockapetris 1987). The DNS is a delegated, distributed lookup system built to enable the real-time translation of host names (such as `www.example.com`) into network addresses, so that clients' hosts can send packets to the desired server machine. The fact that the DNS is highly distributed and delegated is important: at the time of inception, there was no possibility that any single service could provide a globally accessible database to do the lookup in a way that would scale to the number of times that hosts would need to look up addresses, and with the necessary geographic spread. Additionally, because the names are hierarchical, delegation of the management of portions of the domain name space meant that the maintenance (keeping the data up to date) was done closest to the organization that is particularly interested in, and able to provide, accurate information. For example, a Web server manager is in a position to know when the Web server's host name entry in the DNS needs to be updated.

In order to be part of the Internet, all hosts running such application and infrastructure services are expected to abide by the defined standards for the services, and by best practices.

PROPRIETARY SERVICES

As the Internet evolved and spread, a set of specialized and well-known services grew up on and around it. While the WWW (and Gopher³ before it) was intended to be the foundation for collecting and serving managed information sources, it didn't take long for some of those sources to become better known than others (Anklesaria et al. 1993). Amazon, eBay and Facebook are large companies that use their websites (and other network services) in order to connect to their customers and transact business. The website software they use is based on open standards, but the services themselves are commercial, proprietary and private.

There was a period of time when people found a company's website by guessing its domain name ("`www.<trademark>.com`"). Since finding stuff on the Internet is still a key activity, many people directly or indirectly use a search service, such as Google, for that purpose. Google is a large company whose website has

become well known because the company has earned a reputation for providing its service very effectively. Specifics of technology aside, an important difference between the DNS and Google is that the former is an Internet infrastructure service, based on open standards and operated in the best interests of the Internet, and the latter is a proprietary commercial service.

While people originally used their servers' standards-based electronic mail server to send and receive email, it is increasingly common for people to use a commercial email service (such as those provided by Google and Yahoo!). Commercial email services use ISPs to communicate with other email servers to send and receive email; however, the service they are providing is a private one, governed by the agreement with their customers and not by the Internet's standards.

Clearly, proprietary services are key to the Internet's usefulness, but it is important to understand the distinction between infrastructure and proprietary services when it comes to adopting standards, developing accessible features of the Internet and applying regulation appropriately.

NETWORK OF NETWORKS

Above all else, the Internet is a "network of networks." Created in an era when it was infeasible to build a single globe-spanning network, its purpose then was to take existing local networks (typically research labs or campuses) and join them together so that every network host could reach all others. Three key realities emerged from this:

- Local networks are individually built and managed to serve the needs of the users in the lab, enterprise or customer sites.
- These networks are interconnected by virtue of interoperable protocols.
- Global reach is achieved not only by hooking each individual network up to all others, but rather by sharing resources to connect networks that are far apart.

This has meant that the Internet has required a communal effort since its inception, even as it empowered individual networks to be developed and deployed to suit users' needs. It also means that it is very hard to do something to one part of the network and not affect the Internet as a whole.

³ The Gopher protocol was an earlier application designed for distributing, searching and retrieving documents over the Internet. It organized and presented information in hierarchical menus, easily supported by the text-based terminals commonly in use in the late 1980s.

THE UNVARYING CHARACTERISTICS THAT DEFINE THE INTERNET: THE INVARIANTS

In 2012, the Internet Society published a white paper describing characteristics of the Internet that have been stable through its history — “Internet Invariants: What Really Matters” (Internet Society 2012). These are *unchanging* or *invariant* features or supporting conditions. The thesis of the white paper is that these conditions need to be maintained as the Internet continues to evolve. A network that does not have these characteristics is a lesser thing than the Internet as it has been experienced to date.

As it happens, none of the characteristics have to do with specific technologies used to implement the Internet. Any other network, built using completely different protocols, hardware and services, that still demonstrated these characteristics could be equally welcomed and valued. Indeed, the Internet as we know it has undergone many such changes and evolutions — in ways that do not affect these underlying characteristics. While describing what must remain true about the Internet, the invariants offer insight into areas where much change is possible.

As such, these invariants create a framework through which to look at trends, impacts and possible changes to the Internet and its use. How would these forces impact the Internet in terms of its unchanging characteristics?

GLOBAL REACH, INTEGRITY

Global reach, integrity: Any endpoint of the Internet can address any other endpoint, and the information received at one endpoint is as intended by the sender, wherever the receiver connects to the Internet. Implicit in this is the requirement of global, managed addressing and naming services. (Internet Society 2012)

Often quoted as “the end to end principle,” the Internet is known for supporting connectivity between all endpoints. When the Internet was originally developed, every computer was directly connected to it, and it was expected to support all the services of such “host” machines. This was part of the notion of collaborative networking. Host machines would report status, participate in routing, provide services such as “finger,” “talk,” email (receipt and delivery) and file transport protocol (for sharing files).

The beginning of the end for such true global connectivity came along with the realization that IPv4 address space would be insufficient to provide unique addresses to all computers connecting to the Internet. At that point, users’ computers disappeared behind Network Address Translators (NATs) to share a single IP address, NATs were

embedded in “firewalls” that blocked undesired traffic and connections and the common reality became stub networks attached to access networks (for example, from ISPs) attached to the global Internet backbone.

Nonetheless, although it is tricky and sometimes requires expertise to “punch a hole” in your household firewall, it is still generally possible for two computers to connect to each other directly through the global Internet, no matter what networks they are attached to.

The integrity of the Internet extends to its infrastructure services. There have been many discussions of the importance of a single root of the DNS (Internet Architecture Board 2000). The inherent requirement is that one person gets the same view of the Internet (same answers from the DNS) as their neighbour, or someone from across the planet.

Note that there is a subtle difference from ubiquitous proprietary services: DNS is an authoritative Internet infrastructure, designed to provide that uniform view; Google is a proprietary service, which might provide more satisfactory results by tailoring them to different locales. Whether results should be identical across geographies is a business question for Google, not a question of Internet integrity.

GENERAL PURPOSE

General purpose: The Internet is capable of supporting a wide range of demands for its use. While some networks within it may be optimized for certain traffic patterns or expected uses, the technology does not place inherent limitations on the applications or services that make use of it. (Internet Society 2012)

The Internet was not built for any particular application. It was not designed to support a particular activity, such as voice communications or video program delivery. Among other things, this means that there are no a priori assumptions about endpoints or chokepoints or ebb and flow of data on the network. While ISPs are geared toward serving customers, there is no architectural equivalent of “subscriber” in the Internet’s technology. There are the Internet hosts, which are the connected endpoints. Originally, they were fully-fledged server machines and workstations, running a full suite of Internet service programs. Now, they vary from racked multicore data servers to personal computers to hand-held devices and car components. Even so, there is no distinction in Internet network protocols to account for the difference in endpoint type. Indeed, this type of diversity and proliferation of network-enabled devices would not have been possible if there was some finite list of known and supported hardware.

Nor is the Internet multi-faceted, supporting a fixed range of applications and services, which, bundled together, seem like a wide enough array of services to be considered general. Any given device must use standardized networking protocols in order to communicate over the Internet, but the communication of data to support applications and services may be through standard protocols (such as HTTP for the Web, or SMTP and IMAP for sending and retrieving email), which are openly specified and identified in the communicated packets. In keeping with the general purpose nature of the Internet, however, it is to be understood that new protocols will be developed and, therefore, the list of possible protocols is not closed or even finite.

This is not to say that networks cannot be usefully studied and optimized. Rather, optimization has to be at the level of objective measure of packet traffic and not making choices based on endpoint or application type. For example, the Internet Engineering Task Force's (IETF's) Congestion Exposure Working Group is specifying how to signal congestion experienced so that appropriate traffic management decisions can be made. Since the network architecture does not inherently support differentiation between applications, tweaking a network to respond differently to applications based on "deep packet inspection" and "heuristics" (which amount to guesses) derails the generality of the network and its potential uses.

SUPPORTS INNOVATION WITHOUT REQUIRING PERMISSION

Supports innovation without requiring permission (by anyone): Any person or organization can set up a new service, that abides by the existing standards and best practices, and make it available to the rest of the Internet, without requiring special permission. The best example of this is the World Wide Web — which was created by a researcher in Switzerland, who made his software available for others to run, and the rest, as they say, is history. Or, consider Facebook — if there was a business approval board for new Internet services, would it have correctly assessed Facebook's potential and given it a green light? (Internet Society 2012)

It seems reasonably well understood that the open nature of the Internet, as captured in the other invariants, acts as a basis for allowing anyone to make use of the Internet. It is, though, important to remember that "using" the Internet means more than being able to download existing content or connect to services. It also means being able to create and share content, build new services and build new networks/parts of the Internet.

This is not to suggest that there are no rules of the road, or that the Internet is a free-for-all. There are protocols for passing traffic on the Internet, and anything failing to observe those protocols will be ignored or dropped. It does, however, suggest a key distinguishing feature from other large networks, such as the electricity grid and telephone networks, which are both tightly monitored, operated and controlled by government and industry. For good reasons, which are tightly coupled with the approach to development of those networks, it is not the case that anyone can decide to modify their phone's interaction with the telephone network or offer new dialing services on the telephone network itself.

For the Internet, permission-less innovation is not simply an interesting side effect, or a "nice-to-have" feature of the network. The fact that innovation (of the network and of the services that run on it) can come from anywhere has meant that the growth and evolution of the Internet is not limited by the imagination of some collected group of governing minds. The Internet can leverage the creative power of every person in the world. As noted in the description of the invariant, that has brought some unpredictably successful results.

This is not just a historic perspective. School children and hobbyists around the world are building their own special-purpose computing devices based on the Raspberry Pi, a credit-card-sized general purpose computer that supports Ethernet connections.⁴ There is no telling where these devices will turn up or what they will be doing — and that is a good thing, from the standpoint of supporting maximum innovation and evolution.

This approach goes hand in glove with the characteristic that the Internet is a "general purpose network."

ACCESSIBLE

Accessible — it's possible to connect to it, build new parts of it, and study it overall: Anyone can 'get on' the Internet — not just to consume content from others, but also to contribute content on existing services, put up a server (Internet node), and attach new networks. (Internet Society 2012)

As a network of networks, there is no fixed form or function of network architecture. Any one network can connect to one or more other networks, building out the edges of the Internet or creating more interconnection routes. This makes the Internet more than some great wishing well of content into which everyone can dip: anyone can play a more active role than simply accessing existing content and services on the Internet.

⁴ See www.raspberrypi.org.

The heterogeneity of the Internet also lends itself well to study. Any individual can gain insight into network connection status through the use of a few simple command line tools. There is no single or small collection of controlling entities that control “the network,” decide what to monitor in it and, of that, what to publish. Some networks and third parties analyze everything from connections to access speed, via direct analysis and participating probes.⁵ This makes the Internet much more transparent than typical telecommunications networks or electricity grids.

That transparency is advantageous for those looking to improve overall network performance. For example, it is possible to demonstrate the need for, and impact of, IXPs in Africa and elsewhere by demonstrating the before and after impact of installation.

INTEROPERABILITY AND MUTUAL AGREEMENT

Based on interoperability and mutual agreement: The key to enabling inter-networking is to define the context for interoperation — through open standards for the technologies, and mutual agreements between operators of autonomous pieces of the Internet. (Internet Society 2012)

“Interoperation” is the basis of internetworking: allowing separate networks, built with differing hardware, to connect and communicate consistently. This is achieved by having set standards to which equipment must be built and networks set to operate.

Strictly speaking, those standards can be proprietary to a particular corporation or closed consortium of companies. They might be made available freely, or for some price (small or large). They might be made available only to certain authorized parties (for example, certified companies). However, that is not the general model of Internet standards. By ensuring that standards are not only freely available, but also developed through open processes, components of the Internet can be developed by the broadest range of developers. New and different types of networking equipment can be built to connect to the Internet.

“Mutual agreement” is also key to this model of operation. Rather than legislated sets of standards, and regular review thereof, networks participate in the Internet and make connections based on mutual agreement. Standards are voluntarily adopted.

COLLABORATION

Collaboration: Overall, a spirit of collaboration is required — beyond the initial basis of interoperation and bilateral agreements, the best solutions to new issues that arise stem from willing collaboration between stakeholders. These are sometimes competitive business interests, and sometimes different stakeholders altogether (e.g., technology and policy). (Internet Society 2012)

The Internet (internetwork) was created out of a need for collaboration — connecting researchers at disparate centres and sharing resources. While collaboration may be perceived as an obvious form of interaction for research centres, the spirit of collective stewardship of the network and collaboration to fix problems persists in today’s heavily commercial, global Internet.

The IETF was formalized in 1986, while the Internet was still driven by research and academic networking efforts. It adopted a spirit of collaboration to develop technical specifications — participants in IETF discussions are expected to contribute their individual technical expertise and opinion. Successful conclusion of discussion and selection of outcomes is based on determining *consensus* — not voting, not unanimity, but agreement on a majority view.

Collaboration is not limited to the confines of select Internet institutions. Even as the Internet is predominantly made up of commercial networks, operated for profit and in competitive industries, there are times when addressing a larger Internet issue requires those entities to work together in common cause. This was demonstrated very concretely in the World IPv6 Day (June 8, 2011) and World IPv6 Launch (June 6, 2012) events.⁶ With the Internet Society hosting as a neutral party, Google, Yahoo!, Facebook and other content providers — natural competitors — joined forces to demonstrate the feasibility of IPv6 deployment in the face of increasing scarcity of IPv4 addresses. No doubt, there was self-interest involved — Lorenzo Colitti (2009) of Google articulated the need for IPv6 in order to ensure business continuity. But, of the many approaches major content companies could have taken, sharing expertise and contributing to collaborative events is one of the few that demonstrates commitment to the “collective stewardship” framework of managing the Internet.

⁵ See <https://atlas.ripe.net> and www.routeviews.org.

⁶ See www.worldipv6launch.org.

REUSABLE (TECHNOLOGY) BUILDING BLOCKS

Technology — reusable building blocks: Technologies have been built and deployed on the Internet for one purpose, only to be used at a later date to support some other important function. This isn't possible with vertically integrated, closed solutions. And, operational restrictions on the generalized functionality of technologies as originally designed have an impact on their viability as building blocks for future solutions. (Internet Society 2012)

Closely related to the “general purpose” nature of the Internet is the fact that its underlying technologies are created as “building blocks.” Protocols specify what inputs are expected, what outputs will be produced and the conditions on which the former produces the latter.

This building block approach has allowed the Internet to evolve in directions unimagined by its creators. Just as the Internet's routing system does not specify a complete, permanent path (circuit) from one endpoint to another, but leaves it to the routing system to calculate the best path for a packet, technologies get stretched to fit new needs time and time again.

Two key examples are HTTP (the transport protocol for the WWW) and the DNS. HTTP was designed specifically as the communication protocol between Web servers, typically transmitting HyperText Markup Language (HTML) pages of content. With the separation of the definition of the communication protocol from the specification of the content, it was possible to focus on the needs of transmission in defining HTTP. Key things included: establishing credentials and capabilities (of server and client), identifying content being requested and indicating the format of the content being sent. HTTP is tuned to do those things (and other, more detailed actions) very well. At the same time, that's a pretty generic framework for communications of services — whether it is retrieving Web pages or carrying out other services for clients. As a result, HTTP is used for many application services that have nothing to do with strict WWW services. Additionally, it is now common to embed Web servers on special purpose hardware (such as home gateways, microcontrollers and so on), to provide HTML-based configuration tools.

Even before there was HTTP, there was the DNS, set up as a globally distributed lookup service to map domain names to IP addresses. While there is a unique root of the DNS, and it is fundamentally based on hierarchy, another key feature of the DNS is that the detailed information for a domain is maintained under the authority of the domain name holder. Indeed, while it is common to see three-part domain names

today (for example, www.thinkingcat.com), where the domain is essentially a flat list of hosts within the domain (for example, “www”), the DNS can easily be further subdivided in structure and organizational maintenance. For example, www.us.example.com can be maintained and operated by a different administrative group within an Example Company than www.ch.example.com. The expectation is that the administrative staff with the most immediate knowledge of the correct values to store in the DNS will have direct access to the tools to keep it up to date. Put more simply: VeriSign (the registry operator for “.com”) need not update anything in its registry when the administrator of [thinkingcat.com](http://www.thinkingcat.com) moves its website (changing the IP address of www.thinkingcat.com).

Again, taking a step back and looking at the DNS in the abstract, it is tuned as a globally distributed lookup system, keeping the maintenance of current data “closest” to the party responsible for the data. As such, it was straightforward to update DNS to accommodate IPv6 alongside IPv4 — the definition of DNS was not bound to the IP address type of the time. More adventurously, the DNS has been put to different uses — both as a lookup system for things other than obvious domain names (Uniform Resource Names, for example), and to store data (other than IP addresses) associated with domains (Mealling 2002). Some of those other uses of the DNS go to addressing issues that are themselves requirements of the changing nature of the use of the Internet. For example, there are demands for increased security of the Internet's infrastructure, and efforts to reduce unsolicited, and sometimes misleading, email messages (“spam”). Approaches to mitigating those issues require storage of, and access to, so-called “digital security certificates” for ensuring authenticity of the DNS results themselves (see Arends et al. 2005 and related Requests for Comments [RFCs]), or of the authorized mail entities associated with the domain (see Crocker, Hansen and Kucherawy 2011).

Because the DNS is a building block, it is not necessary to establish and deploy a new system for each and every one of these services. Such deployment would be prohibitive for establishing new services.

NO PERMANENT FAVOURITES

There are no permanent favourites: While some technologies, companies and regions have flourished, their continued success depends on continued relevance and utility, not strictly some favoured status. AltaVista emerged as the pre-eminent search service in the 1990's, but has long-since been forgotten. Good ideas are overtaken by better ideas; to hold on to one technology or remove competition from operators is to stand in the way of

the Internet's natural evolution. (Internet Society 2012)

At a technical level, this principle demonstrates how the Internet has continued to evolve to support a wide range of activities that were not conceivable at the time of its inception. Systemically, the Internet supports and fosters approaches that are useful; old, outdated or otherwise outmoded technologies die away.

The same principle applies at the level of use of the Internet — interest in the social networking site MySpace decreased once people determined that Facebook was their platform of choice (see Hartung 2011). Facebook will continue to be the “it” platform until something else comes along that grabs people's attention.

In biological terms, we can say that the Internet supports survival of the population, not the individual. The shuttering of search engine AltaVista did not signal the end of search services for the Internet, just the end of that individual search service. It may have taken with it particular characteristics (traits) that are not found in Google or other search engines, but evolution determined that those were not valuable enough traits to make the service viable.

At the time of this writing, IPv4 is by far the dominant protocol used for Internet traffic, with its successor, IPv6, just beginning to show signs of viable global adoption. Most efforts to promote its uptake have painstakingly emphasized the adoption of IPv6, and avoided the question of turning off IPv4. Networks that “just work” with IPv4 would be threatened by such a prospect. As insurmountable a task as IPv6 deployment is, it would be magnified a thousand-fold if it required the enumeration and treatment of IPv4-dependent networks and devices that cannot migrate (for example, any machine running Microsoft Windows XP, which is long-since past its life expectancy, but still very much in use in odd corners of enterprise networks). At the current rate of adoption of IPv6, which is doubling every year (see the data from Google 2014), IPv6 will be the primary IP used to access Google by mid-2018. Technology pundits who have done the math to rationally predict how long IPv4 will persist as a required network technology suggest it will not disappear altogether before 2148 (that is, over 100 years from now): “At current growth rates, assuming adoption of IPv6 is linear, it will take almost 67 years for IPv6 connections to surpass IPv4 connections and the last IPv4 connection won't be retired until May 10, 2148” (Prince 2013).

An alternative perspective is that IPv4 will, in fact, die away much more rapidly as IPv6 is not only dominant, but also cheaper and easier to maintain. It will become easier to replace IPv4-only systems outright rather than to continue to support them.

Key to all of this is the fact that this process of growth, overtaking existing systems and possibly fading away is quite natural in the Internet. Indeed, it is fundamental to its continued health. It is important not to make policy decisions that in some way lock in a particular technology or implementation. Equally, it is important not to try to prop up businesses or business models that seem to be financial giants. The giants may well fall away — clearing the path for newcomers and an improved Internet.

No only is fighting those trends very difficult, success would mean taking away one of the fundamental drivers of the Internet, and this should be avoided.

SITUATIONAL CHALLENGES AND THREATS OF FRAGMENTATION OF THE INTERNET

This section explores three categories of situational issues that drive different kinds of fragmentation in the Internet. In the first two categories, policies are applied in the interests of making the Internet reflect some level of national agenda. The challenge is how to better achieve that agenda, or resolve the motivation for control, in ways that are more consistent with allowing the Internet to thrive. In the third category, cases where private sector drivers are left ungoverned can create fractions in the Internet.

Each of these challenges is reviewed through the lens of the Internet invariants, to understand how the situation's outcomes can negatively impact the Internet in significant ways. Alternative perspectives are also offered.

ALIGNING THE INTERNET AND ITS RESOURCES WITH NATIONAL BORDERS

This section outlines three cases where there are drivers that would (intentionally or otherwise) put national boundaries on the Internet itself, its resources or its data services. The drivers are based on the rational need to ensure that the Internet and its use are not undermining the fabric of a nation or its citizens' well-being and proper behaviour. However, the approaches taken to make control easier undermine the Internet's integrity, and alternative approaches to international collaboration might provide better avenues for solving the problems.

Putting National Borders on the Internet

The key drivers in this situation are ensuring legal enforcement and control over citizens' actions, and ensuring citizens are not exposed to foreign legal frameworks for inherently domestic activities.

In 2013, revelations of US government data collection practices caused other countries' governments to consider how much of their citizens' traffic flows through the United

States, whether or not it is destined for any user or service there. These realizations have led to calls to reroute major Internet links to avoid having traffic transiting US networks. Changing network connections (and, thus, routes) is a common and ongoing occurrence, but it is usually driven by needs for network efficiency and resiliency. Attempting to re-architect the Internet so that citizens' traffic remains within certain geopolitical boundaries is at odds with responding to the global Internet's needs, and may well lead to less diversity and resiliency in (national) networks.

A look at global connectivity maps provides some surprising information — Internet connections do not naturally align with political boundaries. For example, Canada has an immense geography and a modest population. Population centres (and, therefore, obvious locations for networking hubs) are generally spread apart. Since the Internet's routing technology is designed to pick efficient steps between origin and endpoint, it is not surprising that it is sometimes cheaper, easier and faster to route Internet traffic from one end of Canada to its middle via a connection point in the (much more densely populated) United States, Canada's neighbour to the south. So, traffic from Canadian cities Vancouver to Toronto might reasonably bounce through US cities Seattle and/or Chicago.

Similarly, many international connections out of countries in Latin America terminate in Miami. Miami terminates important data links from other continents. Rather than building individual links between every country in South America to every other continent (or country), it has been most effective and efficient to build large-capacity links to Miami from South America, and have South American traffic transit Miami on the way to or from countries in Europe.

"Cheaper," in the context of interconnections, can mean more than a slight savings for companies involved. However, requiring changes of interconnection to align with country boundaries is more than just a messy and expensive question of network operators changing their connections. It is important in terms of what it means for a resilient, robust Internet.

Through the Lens of the Invariants

Trying to ensure control over citizens' networked life by forcing the Internet's components to line up with national boundaries is directly in conflict with the invariant "global reach, integrity."

The Internet was not designed to recognize national boundaries. It's not being rude — they just weren't relevant. Resiliency...is achieved through diversity of infrastructure. Having multiple connections and different routes between

key points ensures that traffic can 'route around' network problems — nodes that are off the air because of technical, physical, or political interference, for example. We've seen instances where countries are impacted by disaster but at least some of that country's websites remain accessible: if the ccTLD has a mirror outside the impacted network, and if the websites are hosted/mirrored elsewhere, they're still accessible. This can be incredibly important when a natural disaster occurs and there is a need to be able to get to local resources. (Daigle 2013)

Additionally, it is arguable that the more networks align on national boundaries and are perceived as national resources, the harder it is to ensure that the Internet remains "accessible," or that operation must be based on "collaboration," or "based on interoperability and mutual agreement."

Core Policy Perspective

As noted above, the heart of the problem being addressed is nations' desire to ensure their ability to enforce their laws and ensure their citizens are not exposed to foreign legal frameworks for inherently domestic activities. A different approach to ensuring the appropriate treatment of citizens' rights is to work cooperatively to produce effective and enforced laws on appropriate behaviour — on both sides of borders.

Country-based IP Address Allocation

The key driver in this situation is a desire to secure adequate and appropriate Internet resources for one's country, as well as monitoring and/or controlling the management of those resources.

Initially, IP address allocation was a matter of collegial agreement and managed by one person, Jon Postel (see ICANNWiki 2014). With the expectation that the network was destined to connect existing and future research sites, the belief that addresses were plentiful, and the use of hierarchical routing approaches, addresses were handed out in large blocks to single organizations, chiefly in the United States. Those allocations can be seen as "legacy" allocations in the Internet Assigned Number Authority (IANA) registry of IPv4 addresses (see IANA 2014).

Once it became clear that the development of the Internet would outstrip this approach to allocation, the hierarchical approach to allocation and routing was set aside in favour of "Classless" Inter-Domain Routing in 1993 (Fuller et al. 1993). This permitted the allocation of much smaller chunks of IP address space to create usable networks. In the same time frame, the management of allocation of IP addresses

was becoming a task too big for one organization, and the RIR system was established (see more in Karrenberg et al. 2014). Today, there are five RIRs, partitioning the globe, each running open “policy development processes” to develop the allocation and address management policies to apply within region.

With IPv6, addresses are again plentiful. Management in order to control scarcity is not an issue, and with the fresh address space of IPv6, historical imbalances in allocation are no longer relevant. Nonetheless, management of best practices surrounding use and routing are still very timely, and discussions within the RIR open policy development processes are important for ensuring that Internet numbers continue to be used in the best interests of the Internet as a whole.

The careful management of IPv4 address allocation was originally about managing for scarcity, but also for aggregation in inter-domain routing (see Internet Society 2013). That is less of an issue now, with IPv6 and bigger hardware, but the bottom-up, community-driven regional approach is still applicable.

Through the Lens of the Invariants

This is significantly related to aligning operational networks with national borders, and similarly threatens “global reach, integrity.” The pool of IP addresses from which a country would allocate would easily identify that country’s networks, making it easier to prioritize or block entire nations’ networks. It would also move away from the “collaboration” model of RIR open policy development processes, and base allocations on rule of local government rather than focusing on “interoperability and mutual agreement.”

Core Policy Perspective

The problem at hand in this case is that countries wish to ensure they have ample access to appropriate levels of critical Internet resources. Rather than treating resources as a raw material or good that needs to be “owned,” with the attendant impact on the Internet as noted above, countries seeking to ensure that they have appropriate voice in IP address allocation policy going forward could engage in the existing policy process to ensure their concerns are heard and understood. RIR policy discussions are public, and many of the RIRs are performing specific outreach to governments to identify issues and facilitate involvement.⁷

Data Localization

In response to the revelations of government spying, Brazil introduced a proposal in its Internet bill of rights, Marco Civil da Internet, to require global Internet companies

such as Google to establish data repositories within Brazil (Government of Brazil 2011). Although the specific proposal has been dropped from the now-adopted Marco Civil (see Boadle 2014), the concerns that drove it remain. Those concerns are that citizens’ communications are being subject to scrutiny by another nation’s government.

At a distance, it seems perfectly straightforward to assert that users’ communication with large global companies should be carried out uniquely within a user’s country. Expressing that in terms of Internet infrastructure leads to the requirement that data centres be housed in that country.

However, such requirements, if imposed, could easily fall into the category of both failing to achieve the policy objective and stifling the Internet. As an added issue, such requirements may impact users’ experience of the service.

Requiring data centres to be in-country ensures that a citizen’s communications with the service stays within the boundaries of the country if (and only if) the network path from the user to the data centre remains within the country. Unless there are national boundaries on the Internet, or the large corporation is directly serving each access provider (home and business), there are no such guarantees. Additionally, citizens travel, and it is inevitable that some citizens’ interactions will be made through data centres elsewhere in the world.

The user’s experience of connection performance can easily degrade if they are in a remote part of Country A, closer by geography (or, at least, network topology) to a population centre of Country B, where a data centre might reasonably be located. Sizing data centres to meet the needs of each country’s population, with no possibility of failover or offloading⁸ to other data centres is a challenge, which is likely to leave less interesting markets underserved by the corporation.

Through the Lens of the Invariants

This general approach is stifling to the Internet because it undermines its “general purpose” nature (since networks and services are architected to predict and match user transactions), and the “global reach and integrity” of applications. Historically, the focus of service build-out has been on offering resiliency through redundancy and replication, leveraging availability of different networks to provide robustness.⁹ Requiring localized data for large

⁸ Failover occurs when one server cannot continue and a backup server is put into use (seamlessly, it is hoped). Offloading refers to sharing, among several servers, the load of responding to incoming requests.

⁹ For example, although there are still only 13 distinct DNS root servers, many instances of them are now multicast to enable reliable access in all parts of the world, and thus from all over the globe.

⁷ See www.ripe.net/ripe/meetings/roundtable.

services changes the emphasis to focus on consumers' geographic locations.

This approach also threatens the expectation of “innovation without requiring permission,” and “no permanent favourites”: What nascent company can immediately provide separate services in every country on the planet? Or, must services that cannot comply with such requirements block access to would-be users from those countries requiring data localization? In either case, the Internet is impoverished and/or fragmented.

Core Policy Perspective

The issue being addressed is the exposure of citizens' information (Internet usage, transactions, personal information and so on) to companies operating under other countries' laws. An alternative is to look at the issue of data privacy outside the narrow scope of eavesdropping, to develop and enforce policies for the appropriate handling of data. “Appropriate handling” ranges from confidentiality (in transmissions and storage) to conditions under which personal data may or may not be shared. These are not easy issues to address, but addressing them is inevitable, for the sake of the world's societies, if not for the Internet's future.

CONTROLLING ACCESS THROUGH INFRASTRUCTURE RESTRICTIONS

The greatest thing about the Internet is that it erases borders and distance. The most challenging thing about the Internet is that it erases borders and distance. Governments seeking to regulate behaviour in their jurisdictions are often faced with the reality that an activity that is deemed inappropriate is happening *outside* their jurisdiction. Absent international agreement, they have no means to address the issue where it is happening.

Tweaking Local Infrastructure

As a proxy for actual control, governments have on occasion imposed restrictions on Internet infrastructure that is resident within their jurisdictions, instead of aiming to control access to, or engagement in, the offensive activity.

For example, Russia is routinely on Hollywood's watch list of countries not adequately policing piracy of American-made movies (see Block 2014). For many years, servers in Russia have offered unauthorized copies of movies with relative impunity from Russian law enforcement agencies, although enforcement is said to be becoming tougher (see Kozlov 2014). Since all of this is hosted within Russia, there is nothing that US officials can do about enforcement of US laws that prohibit such serving of copyrighted material.

In many ways, this is not a new problem — copies of films have been smuggled out of one country to be viewed in other countries for as long as there has been a movie

industry. However, that has physical limits, and a key difference with the Internet is that the viewers do not have to be in Russia. American viewers can watch a Hollywood movie obtained from a Russian piracy site, as long as they know where the servers are and how to navigate their indexes.

The above illustrates one case of a situation where the government of a jurisdiction believes that inappropriate (illegal or otherwise problematic) services are being offered on the Internet, hosted in another country. A typical, but largely ineffectual, approach to addressing their citizens' access to the services is to curtail Internet access from the home country. In that light, the proposed “Stop Online Piracy Act” (SOPA) and “Protect IP Act” (PIPA) that US senators proposed to control US ISPs' DNS responses to customers, the blockage of DNS resolution for Twitter and YouTube during the 2014 unrest in Turkey (see Letsch and Rushe 2014) and Egypt's outright unplugging of the Internet in 2011 (see Al Jazeera 2011) are all the same. The motivations may be different, but each action seeks to curtail access by controlling (and, in so doing, breaking) local Internet infrastructure.

A slightly different issue occurs when one country acts to prevent anyone from accessing content or services that it deems inappropriate. The US Immigration and Customs Enforcement (ICE) agency has, since June 2010, pursued a program of seizing domain names of sites deemed to be “illegally selling counterfeit merchandise online to unsuspecting consumers” (see ICE 2013). In recent years, ICE has teamed up with related agencies in other countries to broaden the scope of seizures (see EUROPOL 2013). In all cases, law enforcement agencies can only seize domains that are registered with registries housed within their jurisdiction — such as .com, .net and .org, which are operated by companies based in the United States. Typically, these seizures are done because the website hosting the trademark-infringing material is hosted elsewhere (outside the reach of the concerned law enforcement agencies). Once the domain name is seized, ICE trades off the domain name's mark by directing it to ICE's own servers and displaying its own message (on anti-counterfeiting).

Additionally, sometimes there are unintended consequences, such as when Pakistani authorities demanded that YouTube be censored within Pakistan. Pakistan Telecom was (necessarily) responsive, and on February 24, 2008, Pakistan Telecom's routers announced a more specific (appealing) route to YouTube's servers. The intention was to use this to direct Pakistani traffic away from YouTube. Unfortunately, the routing information was not contained within Pakistani networks and was duly propagated through the global routing system — drawing all YouTube traffic to Pakistan Telecom's network and thereby effectively knocking YouTube off the Internet for everyone.

Through the Lens of the Invariants

In all the cases outlined above, the “global reach and integrity” of the Internet and its core services is threatened, leading to fragmentation and disintegration through local exceptions to how the Internet behaves.

Additionally, these approaches undermine the reusable building blocks of the Internet, such as DNS. The SOPA/PIPA proposed legislation made requirements on the use of the DNS for systems. That would curtail the use of DNS going forward, in some ways freezing its current existence as the state forevermore. Put slightly differently, it would reduce its use as a building block technology as if some of the corners had been sawed off the blocks themselves. As noted in the description of the “reusable (technology) building blocks” invariant, there are ongoing technology developments that leverage the DNS infrastructure, and they would be impacted.

More subtly, these approaches undermine the “collaboration” and “mutual agreement” approaches to developing and operating the Internet, because they emphasize that operators are responsive to laws and regulations, not collaboratively building the Internet.

Core Policy Perspective

At the heart of the matter, the objectionable behaviour is occurring outside the jurisdiction of the complaint and thus outside the reach of local (national) laws. However, the Internet and its infrastructure are not the problems in these cases. Instead, effective and enforced laws on appropriate behaviour — on both sides of border — are required in order to address the situations outlined.

DIVERGENT REALITIES BASED ON BUSINESS MODELS

As the Internet is increasingly made up of commercial networks, one of the key ways to influence its evolution, for good or ill, is to focus on the business of building and using it. It becomes important to understand how business decisions and the Internet play together; developing policies for business practices that are supportive of, rather than impediments to, the Internet is key to its ongoing success.

The Internet started as a research network, and was not constructed based on a business model of trying to earn financial profit from operating part of the network or offering services to support it. It has grown to its current scale because compatible business models were found to foster its commercial growth. As a side effect of being (primarily) composed of commercial networks, carrying traffic for commercial interests, business models drive much of today’s Internet shape.

In the general scheme of things, this keeps a healthy balance on deployment of practical advances. Network operators are in the best position to understand how traffic flows through their networks and how to support its use effectively and efficiently. Sometimes, however, necessary services or advances are not well aligned with individual business models, or require a perspective that spans more than the reach of one business’s network in the Internet.

Internet-wide Services

As part of the original Internet set up, several information services were maintained and operated on behalf of the entire network. Network Time Protocol (NTP) is one such service, providing clock synchronization for all interested hosts on the network. The service is a relatively lightweight task and today almost 4,000 NTP servers are available and accessible publicly.¹⁰

As noted above, the DNS was established as another such infrastructure system. Apart from the 13 independent root servers, which provide up-to-date information on finding the so-called top-level domain (TLD) name servers, the initial TLD services were originally defined in memo RFC0920 (Postel and Reynolds 1984), and operated by (or for) the United States Defense Advance Research Agency. DNS is critical to virtually every Internet transaction. Openness and uniformity of the Internet are based on the expectation that every host is equally accessible — domain names are just strings of characters to the Internet’s technology, and anything that made one preferential over another, or impeded access to them, would be harmful to that openness.

And yet, providing domain name service at the TLD level cannot be called a “lightweight” task. Generic TLD registry receives a fixed fee for every domain name registered in the TLD, whether it is for an obscure site or one that is used by millions of people every day. Registries are obliged to scale their services based on resolution demand, which may or may not grow sympathetically with the number of domain names registered in the registry (revenue). In the old telephony model, companies billed a miniscule charge “per dip” into their number database to look up a phone number. Although each charge was miniscule, it added up to revenue. Domain name registries familiar with this model might expect compensation for each DNS lookup, whether from the entity looking up the domain name or the registrant of the popular domain name. However, this is exactly the kind of preferential treatment/impediment to access that is antithetical to the Internet’s success. The fact that no such “per dip” charge has been implemented by TLD operators is key to the Internet’s continued success.

However, this lack of obvious funding model for serving the DNS has perhaps created a resistance to deploying new

¹⁰ See www.pool.ntp.org for details.

Internet-wide services, such as “identity management” providers, or even separate lookup and resolution services for cryptography certificates. Instead, more systems look to leverage the existing DNS infrastructure rather than motivating deployment of another global infrastructure.

Through the Lens of the Invariants

Requiring a business case in order to deploy new technology and services does undermine the “general purpose” nature of the Internet: to the extent that new things must be offered as (private) services, the general purpose nature does not evolve.

Additionally, to the extent that new services are offered on a strictly commercial (and often proprietary) basis, they are not particularly “accessible.”

Core Policy Perspective

The challenge discussed here is that the Internet relies on core services that are offered neutrally and openly across the Internet, where the operation itself bears a cost that is not insignificant. There is relatively little to address this from a policy perspective, except perhaps to provide support for infrastructure services on a public service basis.

Deploying Global Infrastructure Updates

Even as network operators the world over acknowledged that IPv4 address space was running out, it has been very difficult to motivate deployment of equipment and software to support IPv4’s successor, IPv6. That is, although network engineers can articulate the technical impossibilities of running networks without new IPv4 addresses, and the ease with which the Internet can continue to function as a global network once IPv6 is deployed, IPv6 deployment started about 15 years later than intended. At least in part, this is because support for making those investments was blocked on senior executives’ desks for the better part of a decade. The sticking point was that deploying IPv6 was an expense without any perceived near- or medium-term revenue advantage. Indeed, there was little advantage to deploying IPv6 unless or until many other networks and content sources implemented it. This equation changed thanks to the collaboration of several network operators and content companies that worked together to demonstrate the value of breaking the chicken and egg problem, leading the way with significant IPv6 deployment and traffic after World IPv6 Launch in 2012.¹¹

Through the Lens of the Invariants

In order to ensure the “global reach and integrity” of the Internet, it is important to press on with deployment of IPv6 to the point of rendering IPv4 obsolete and unused globally. But IP addresses are not the only needed technology upgrade. A technology designed to address key shortcomings in the level of security of the DNS, DNS Security Extensions, has similarly faced an uphill battle for deployment. Changes to the underlying transmission layer of the Internet are all but impossible because of the need for universal uptake for the sake of compatibility and/or in order to deliver on performance improvements. In any of these cases, partial deployment of a technology infrastructure improvement can lead to fragmentation of the Internet.

Similarly, infrastructure improvements that are achieved by single companies deploying proprietary systems can lead to less “interoperability and mutual agreement” and create monopolies that defy the invariant property of the Internet having “no permanent favourites.”

Core Policy Perspective

The issue being identified is that the Internet does need periodic updating of its core operations, for the good of the Internet as a whole (but not necessarily immediately, or uniquely, for the good of the network operator). Different countries tried varying policy approaches to mandate or encourage IPv6 deployment, with inconsistent levels of success. Generally, policy approaches that foster competition and encourage ongoing upgrading of infrastructure are appropriate.

Charging Models

In 2012, the European Telecommunications Network Operator’s (ETNO’s) association submitted a proposal (ETNO 2012) to the Council Working Group preparing the International Telecommunications Union treaty-developing World Conference on International Telecommunications. The proposed text became known as the “sender pays” proposal for changing Internet business models. Like the load on the DNS registry servers, access networks must scale to meet the needs not only of data sent by their customers, but also data sent toward their customers, chiefly by content providers. The premise of the proposal is that the access networks have no share of the revenue that traffic provides the content distributors, even as the cost of delivery is on the access network. The details of the proposal are not material, insofar as it was just one representative instance of the kind of business logic that has surfaced before and will come to light again. The heart of the issue is that, again, such an approach would throw up roadblocks to the Internet’s flat, non-discriminatory nature. Not all services would be made available across

¹¹ See www.worldipv6launch.org/.

all access networks, and a different form of fragmentation would occur.

Through the Lens of the Invariants

Changing charging models for the Internet to focus on the business overlays (rather than the network interconnections and general carriage of traffic) could have serious impacts on the “global reach and integrity” of the Internet as noted above.

It could also impact “innovation without permission,” insofar as the charging model makes new services prohibitively expensive to new entrants, thereby undermining “no permanent favourites.”

It is completely at odds with the expectation of “collaboration.”

Core Policy Perspective

The claim at the centre of this proposal was that the Internet needs a different business model. From a policy perspective, the best approaches to address the discussion and avoid the negative outcomes of overrunning the invariants is to ensure appropriate anti-competition laws are in place, and to ensure that the Internet remains open to all legitimate traffic indiscriminately.

CONSIDERING THE NATURE OF THE INTERNET IN POLICY DISCUSSIONS

TEASING ISSUES APART TO FIND “WHAT” THE PROBLEM IS NOT “HOW” TO SOLVE IT

The previous section outlined situational challenges for which proposed and existing solutions are at odds with the Internet’s invariant properties: current course and speed may lead to fragmentation of the Internet. Nevertheless, the issues are real and accompanied by a sense that something needs to be done. Each section concludes with a focus on the heart of the problem being addressed, independently of the Internet.

Generally speaking, when there have been issues with the Internet or its use, changes have followed to address the problem. When the source of the issue is behaviour that is external to the Internet itself, forcing change on the Internet typically leads to fragmentation and damage. Therefore, focusing on what the problem is — difficult though it may be — is the best path to follow in order not to undermine the Internet. This often requires stepping back and focusing again on the actual outcome or behaviour that is in question, not the Internet technology that may be involved.

DOES THE PROBLEM NEED A POLICY SOLUTION?

When it comes to considering policy options, the nature of policy needs to be weighed in the light of that fluidity. Policies, laws and international treaties are carefully crafted in the moment and intended to apply for the long term. Volatility is not desirable in policy frameworks — changing them can be long, costly and difficult. The last two decades of the Internet’s history have seen it driven by (largely) private companies’ agreements and efforts. Business agreements are established and torn down relatively easily and frequently. It might be expensive, but costs are factored into decisions to establish and dissolve business agreements. In fact, many business agreements include conditions for dissolution and explicit agreement as to how to wind up the agreement from the outset.

While both laws and business agreements are written to fit the purpose of a given moment in history, the very persistent nature of laws causes them, and regulatory policy derived from them, to freeze the moment in time. They need to be based on what is right and real for the long term; otherwise, they run the risk of making a transient situation permanent. This can be problematic in the long run, in that the future may not be best served by that vision of the Internet.

As a global platform, the Internet has truly thrived since the private sector took on operation of access and transit networks in the 1990s. Not only does the topology of the network look very different today, the technologies and systems running it have evolved commensurately to accommodate greater traffic, and new traffic flows, patterns and network uses.

A CASE HISTORY: PEERING

These growth patterns are not without criticism. “Peering agreements” — business arrangements whereby operators of networks agree to pass traffic for payment or other considerations, have long been the subject of calls for greater transparency and regulation. There is a basic question of level of fairness or competition that is allowed by an industry based on private peering.

If legislation had been put into place in the 1990s to address this and/or enforce outcomes for peering agreements, the landscape of the Internet would have been different — the flipside of open competition is the ability to build business. At the same time, private peering agreements where top-tier companies have a stranglehold on the industry create the kind of “immortal” top dogs that go against the invariant of “no permanent favourites.” Private peering agreements were not the right answer for the Internet, nor was regulation capturing the status quo and controlling it. What we have seen in the intervening decades is the development of other means of Internet

information exchange (specifically, public peering [IXPs], other collaborative arrangements and the build-out of much larger spans of networks). Not only has the industry largely coped with the worst of the competition issues, it has done so by building out new connection arrangements that are more suited to the Internet of today than the simple peering agreements of yore — which would have become entrenched reality with ill-suited legislation.

That said, there are real issues of impact if companies de-peer — for example, in 2008, ISPs Cogent and Sprint had a business disagreement that led to Sprint de-peering Cogent. The consequence of that network change was that uninvolved customers of the two companies were left unable to communicate directly over the Internet (Ricknäs 2008). One question is whether it is appropriate for companies to take an action knowing that it will have that kind of impact on Internet users. However, that's not a question of peering, per se.

FOCUSED POLICY APPLICATION

Policy is set when there is behaviour or an outcome that is desired or should be prevented. In the case of peering arrangements, there may be a desire to “level the playing field” for some competitive interests, or to prevent companies’ business choice implementations from knocking out Internet access for unsuspecting (and uninvolved) users. In the case of the proposed SOPA/PIPA legislation, the outcome that was to be prevented was US citizens’ access to sites accused of online copyright infringement and online trafficking in counterfeit goods.

The challenge, in the latter case, is that the outcome is very hard to prevent or police and the enforcement of laws governing behaviour is difficult. The next logical step, therefore, was to look at the mechanisms that enable the undesired outcome, and curtail the use of them. It is generally easier to control and impose restrictions on computers, software and networks than humans. But, as noted earlier, restricting the technology is poor imitation of achieving the desired goal, because it is so ineffective and has significant collateral damage — to the Internet as it stands today, and to any future growth (of the Internet technology’s building blocks).

CONCLUSION

The Internet is no accident, and while it has developed through evolution in response to changing requirements, its development has not been random or without thought. There are key properties of the Internet that must be supported in order for it to enjoy continued success.

It is no longer possible to grasp the nature of the Internet without considering the world in which it exists — as such, technology considerations may be at the heart of determining what works (or doesn’t) for the Internet, but

a non-technical framework for discussing eventual trade-offs is imperative.

The invariants can serve as a useful framework for discussing impacts without having to delve into the intricate details of the technology that drives the Internet. With the framework in mind, policy discussions can focus on what can be done to address an issue and evaluate potential impacts on the Internet.

WORKS CITED

- Arends, R., R. Austein, M. Larson, D. Massey and S. Rose. 2005. “DNS Security Introduction and Requirements.” RFC4033, March. www.rfc-editor.org/rfc/rfc4033.txt.
- Al Jazeera. 2011. “When Egypt Turned Off the Internet.” January 28. www.aljazeera.com/news/middleeast/2011/01/2011128796164380.html.
- Anklesaria, F., M. McCahill, P. Lindner, D. Johnson, D. Torrey and B. Albert. 1993. “The Internet Gopher Protocol (A Distributed Document Search and Retrieval Protocol.” RFC1436, March. www.rfc-editor.org/rfc/rfc1436.txt.
- Block, Alex Ben. 2014. “India Joins China, Russia, Switzerland on Privacy Watch List.” *Hollywood Reporter*, June 24. www.hollywoodreporter.com/news/india-joins-china-russia-switzerland-714572.
- Boadle, Anthony. 2014. “Brazil to Drop Local Data Storage Rule in Internet Bill.” Reuters, March 18. www.reuters.com/article/2014/03/19/us-brazil-internet-idUSBREA2I03O20140319.
- Colitti, Lorenzo. 2009. “IPv6: The Acceptance Phase.” Presentation given at IETF74, March 22–27. www.isoc.org/isoc/conferences/ipv6panel/docs/colitti.pdf.
- Crispin, M. 2003. “Internet Message Access Protocol — Version 4rev1.” RFC3501, March. www.ietf.org/rfc/rfc3501.
- Crocker, D., T. Hansen and M. Kucherawy, eds. 2011. “DomainKeys Identified Mail (DKIM) Signatures.” RFC6376, September. www.rfc-editor.org/rfc/rfc6376.txt.
- Daigle, Leslie. 2013. “Provoking National Boundaries on the Internet? A Chilling Thought...” Internet Society, June 17. www.internetsociety.org/blog/2013/06/provoking-national-boundaries-internet-chilling-thought.
- ETNO. 2012. “CWG-WCIT12 Contribution 109.”
- EUROPOL 2013. “690 Internet Domain Names Seize Because of Fraudulent Practices.” December 2. www.europol.europa.eu/content/690-internet-domain-names-seized-because-fraudulent-practices.

- Fielding, R. and J. Reschke. 2014. "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing." RFC7230, June. www.rfc-editor.org/rfc/rfc7230.txt and related RFCs.
- Fuller, V., T. Li, J. Yu and K. Varadhan. 1993. "Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy." RFC1519, September. www.ietf.org/rfc/rfc1519.txt.
- Goble, Gordon. 2012. "Top Ten Bad Tech Predictions," *Digital Trends*, November 4. www.digitaltrends.com/features/top-10-bad-tech-predictions/8/.
- Google. 2014. "Google IPv6 Traffic." www.google.com/intl/en/ipv6/statistics.html.
- Government of Brazil. 2011. "Marco Civil proposal of November 2011." [In Portuguese.] http://edemocracia.camara.gov.br/documents/679637/679667/Marco+Civil+da+Internet++6_11_2013/0e3fae49-7e45-4080-9e48-c172ba5f9105.
- Hartung, Adam. 2011. "How Facebook Beat MySpace." *Forbes*, January 14. www.forbes.com/sites/adamhartung/2011/01/14/why-facebook-beat-myspace/.
- IANA. 2014. "IANA IPv4 Address Space Registry." www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml.
- ICANNWiki. 2014. "Jon Postel." http://icannwiki.com/Jon_Postel.
- ICE. 2013. "ICE, International Law Enforcement Seize 706 Domain Names Selling Counterfeit Merchandise." December 2. www.ice.gov/news/releases/ice-international-law-enforcement-agencies-seize-706-domain-names-selling-counterfeit.
- Internet Architecture Board. 2000. "IAB Technical Comment on the Unique DNS Root." RFC2826, May. www.ietf.org/rfc/rfc2826.txt.
- Internet Society. 2012. "Internet Invariants: What Really Matters." www.internetsociety.org/internet-invariants-what-really-matters.
- . 2013. "A Fine Balance: Internet Number Resource Distribution and De-centralisation." www.internetsociety.org/fine-balance-internet-number-resource-distribution-and-de-centralisation.
- Karrenberg, Daniel, Gerard Ross, Paul Wilson and Leslie Nobile. 2014. "Development of the Regional Internet Registry System." *Internet Protocol Journal* 4, no. 4. www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_4-4/regional_internet_registries.html.
- Klensin, J. 2008. "Simple Mail Transfer Protocol." RFC5321, October. www.ietf.org/rfc/rfc5321.txt.
- Kozlov, Vladimir. 2014. "Russia's Anti-Piracy Law to be Toughened, Producing Exec Says" *Hollywood Reporter*, October 7. www.hollywoodreporter.com/news/russias-anti-piracy-law-be-738693.
- Letsch, Constanze and Dominic Rushe. 2014. "Turkey Blocks YouTube amid 'National Security' Concerns." *The Guardian*, March 27. www.theguardian.com/world/2014/mar/27/google-youtube-ban-turkey-erdogan.
- Mealling, M. 2002. "Dynamic Delegation Discovery System (DDDS) — Part Three: The Domain Name System (DNS) Database." RFC3403, October. www.ietf.org/rfc/rfc3403.txt.
- Mockapetris, P. V. 1987. "Domain Names — Concepts and Facilities." RFC1034, November. www.rfc-editor.org/rfc/rfc1034.txt and updates.
- Postel, J. and J. Reynolds. 1984. "Domain Requirements." RFC0920, October.
- Prince, Matthew. 2013. "Happy IPv6 Day: Usage On the Rise, Attacks Too." *CloudFlare* (blog), June 6. <http://blog.cloudflare.com/ipv6-day-usage-attacks-rise>.
- Rekhter, Y., T. Li and S. Hares, eds. 2006. "A Border Gateway Protocol 4 (BGP-4)." January. www.ietf.org/rfc/rfc4271.txt.
- Ricknäs, M. 2008. "Sprint-Cogent Dispute Puts Small Rip in Fabric of Internet." *PC World*, October 31.

ABOUT THE AUTHOR

Leslie Daigle has been actively involved in shaping the Internet's practical evolution for more than 20 years. She was an appointed member of the Internet Architecture Board for eight years, and elected as its chair for five of those years.

Leslie was most recently the Internet Society's first Chief Internet Technology Officer. She helped to (re)create the global dialogue on important technical issues, calling stakeholders to action by providing achievable targets and facilitating their own collaboration across (corporate) organizational boundaries until May 2014.

She is currently principal at ThinkingCat Enterprises, where she has launched the online InternetImpossible.org storybook of the Internet's experienced global impact.

**SECTION TWO:
THE ECONOMICS OF OPENNESS
AND FRAGMENTATION**

**CHAPTER TWO:
ADDRESSING THE IMPACT OF DATA LOCATION REGULATION
IN FINANCIAL SERVICES**

James M. Kaplan and Kayvaun Rowshankish

Copyright © 2015 by McKinsey & Company

INTRODUCTION

A free, global and open Internet has driven tremendous innovation and provided enormous value over the past decade. Financial institutions have used the Internet to establish private communication networks that effectively support their heavy volume of online transactions. This is particularly true of larger firms, which are pushing to globalize their operating models and technology platforms. Global platforms connected by private networks allow banks to provide their best products in all markets, manage risks globally, innovate efficiently and ensure a seamless experience for customers. The advent of sophisticated mobile platforms promises to make all sorts of household and corporate financial management even smoother and more intuitive. Consumers can check balances, make payments and oversee their investment portfolios. Corporations can manage cash positions around the world using increasingly sophisticated online tools.

Recently, countries in every part of the world have issued rules about how corporations must handle their customers' data, including its transmission across borders through the Internet and private networks. Typical motivations for these rules are prevention of cybercrime, protection of citizens' privacy and promotion of the local economy by enforcing job creation. Given recent concerns about data security and privacy, implementation of these regulations has accelerated. In some cases, countries have started to discuss creating their own internets — an emerging phenomenon often referred to as the “splinternet”¹ — through which they would have much more control over the nature of online transactions.

The implications of this new wave of data location regulation are particularly significant for banking. The financial industry has historically been heavily regulated in many ways, such as products offered or capital requirements. As more banking activities, such as data management, come under regulatory scrutiny, the effects are uncertain, but seem likely to be material. In order to provide a context for discussions of governance of the Internet and cyberspace more broadly, executives at a dozen global financial institutions were interviewed and asked about how this complex regulatory environment is affecting financial organizations.

The interviews revealed several implications for banks and, more broadly, markets for financial services. Increasing data location regulations may cause banks to exit some markets, leaving customers in those countries with reduced options for financial services. Those banks

that choose to stay can mitigate the impact of data location regulations with investments that make their technology platforms more modular and flexible. For policy makers, understanding the impact on consumers and, where possible, mandating outcomes rather than specific technology configurations, can avoid or limit any unintended consequences for consumers' access to financial services.

THE VARIETY OF REGULATIONS IS WIDE AND COMPLEX

Countries are creating a wide variety of data location requirements that impose restrictions on the content that has traditionally been transmitted through the Internet (or private networks). These have different implications for the ways that financial institutions manage data. Executives highlighted four main categories of emerging regulations, from most to least stringent:

- **Geographical restrictions on data export**, which require data to be stored and processed within the country (i.e., “data copy cannot leave”). This can force institutions to create separate infrastructure, computing capabilities and teams. Examples include South Korea, which prohibits the export of customer data, and Egypt, which requires banks to keep all information on their government customers within the country.
- **Geographical restrictions on data location**, which allow data to be copied outside of the country for processing, but require a replica in the local infrastructure (i.e., “data copy must stay”). These are, in most cases, motivated by an intention to develop the internal economy. Indonesia and Malaysia are among the countries that do this.
- **Permission-based regulations**, which require institutions to gain consent from individuals for data transmission. For example, Brazil and Argentina require banks to get a customer's explicit written approval to transfer their data. Switzerland and Luxembourg empower customers to prohibit banks from sending their data across the border.
- **Standards-based regulations**, which allow institutions to move data freely outside of the jurisdiction, but require them to take steps to ensure the security and privacy of customer data.

Independently of these levels of stringency, countries can have very different levels of coherence and clarity in their regulatory regimes.

On one hand, these regulations are almost all national rules; as such, they are highly variable — and even contradictory — between jurisdictions. In some cases,

¹ Splinternet is defined as “a characterization of the Internet as splintering and dividing due to various factors, such as technology, commerce, politics, nationalism, religion, and interests.” See <http://en.wikipedia.org/wiki/Splinternet>.

multiple jurisdictions may govern the same data set, and it may be impossible to comply with all mandates. For instance, the United States has protocols on anti-money laundering (AML) and suspicious activity reporting (SAR). To succeed, the protocols must be applied globally, but data location regulations in many countries hinder the necessary exchange of information. Regulations such as the Republic of Korea's privacy requirements, Spain's Data Protection Law and even the United States' own Gramm-Leach-Bliley Act and Non-Public Personal Information Act stand in the way of successful application of AML and SAR protocols.

On the other hand, executives reported that they have severe difficulties gaining a clear and comprehensive view of the full set of regulations. Many are worded so vaguely that it is impossible, they say, to predict what is and is not allowable. In some countries, regulators have given different answers to different institutions, making it difficult to find relevant precedents. Among the concerns they expressed were the following:

- A country's regulations can be worded vaguely (for example, no clear definition of some key terms).
- Some countries lack explicit rules for banks to seek approval of offshore support, leading executives to believe that institutions in the same circumstances receive different treatment.

In fairness, the rules in many countries are still under development. But that only adds to the problem: the uncertain environment makes it particularly difficult to plan and execute large technology investments. Some countries require a number of approvals for an individual compliance project, making it difficult for institutions to plan ahead, and risking delays and significant sunk costs if approvals are not forthcoming.

Institutions will need to live for a decade or more with the implications of the data architecture decisions that they make today. That fact is fundamentally disconnected from the uncertain and variable regulatory rules they now face. Moreover, new technology models may outpace regulations put in place just a few years ago. For instance, several executives said that there is a basic disconnect between cloud computing and the regulatory frameworks in many countries.

EFFECTS ON BUSINESS FALL ALONG A CONTINUUM

Almost all the executives interviewed said that data location requirements are complicating their long-standing strategies to consolidate technology platforms and business operations on a regional, if not global, basis. Otherwise, though, the requirements affect institutions

and countries differently, with the impact falling along a continuum (listed below from low to high impact).


- **Increased organizational complexity to manage.** The complexity of dealing with data location regulations adds another challenge for managers to overcome. However, this complexity does not fundamentally alter business economics at a country or enterprise level.
- **Lower efficiency.** Data location regulations reduce efficiency by requiring institutions to retain people and technology in local markets that they otherwise would not require, reducing margins and resources available for reinvestment. The rules degrade a financial institution's ability to provide service in a seamless way to customers across countries and regions. Banks doing business outside their home countries can face significantly higher costs than domestic banks (see Figure 1), in part because of local data location regulations and in part because of other causes, such as their typically smaller scale than local players. The role played by data location regulations is significant; a recent case study shows that the efficiency ratio of one bank's foreign subsidiary fell by three percent when the bank had to create local data infrastructure (see Figure 2). That was a significant blow, given the industry's competitive margins in this country.

The debate over net neutrality began in the late 1990s in the United States. Since then, it has been gaining momentum in several fields, generating dichotomous positions between different sectors. As a contribution to the debate, this chapter attempts to separate the unquestionable principles — such as the need to preserve the Internet as a space that is open to innovation, and the freedom of users to access content and services — from the dogmas and beliefs that are put forward in the name of neutrality, but which affect the sustainable development of the digital ecosystem.

Telecommunications networks and services and providers of content over the Net uphold the digital ecosystem, and it is essential that both can develop sustainably, with equivalent regulations and principles. This raises two important thoughts. First, it is important to promote investment, innovation and competition, preventing distortions through the relationships produced within the digital ecosystem. Second, the regulatory principles should be balanced between the different actors of the value chain. Meeting certain basic principles in favour of competition and against arbitrary discrimination would create the conditions for fostering the development of the digital ecosystem.

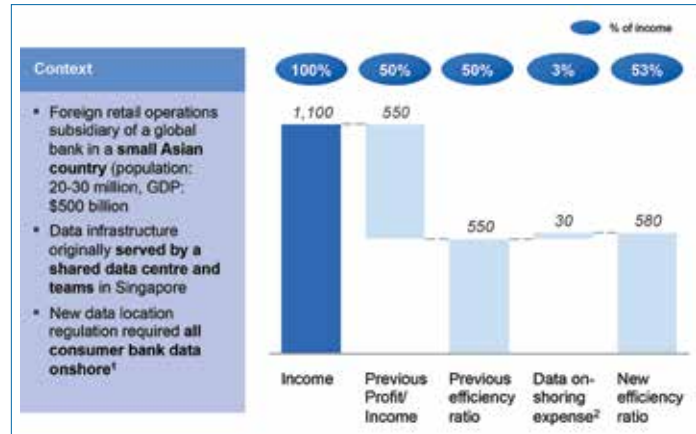
- **Reduction of the global footprint.** Data location regulations make some countries economically unattractive, causing institutions to exit, and

Figure 1: 2013 Efficiency Ratios
(Cost Incurred to Produce \$1 of Income, %)



Data sources: Turkstat, BKM, MasterIndex report, SIS, Central Bank of Russia, SNL, bank annual reports, ViewsWire, Austin Asis, McKinsey & Company analysis, Central Bank (Egypt), World Bank, CIA, UN database, World Bank financial inclusion database, FSS Korea.

Figure 2: Impact of Data Location Regulations
(\$ Millions Efficiency Ratio)



Source: Case example.

Notes: 1. Only production environments; 2. Average of costs in a seven-year period; includes annualized investments and recurring costs.

limiting their global footprint. One bank is already in the process of exiting two countries. Another is considering exiting a country because staying in business there would require tens of millions of dollars in data centre investments.

- **Reduced access to financial services.** As banks reduce their operations due to an unfavourable environment, the expansion of financial services in those countries will slow. This is particularly concerning as the countries that are adopting a more stringent perspective in regulation have the most need to foster development (see Figure 3).
- **Challenges to global technology strategies.** Data location regulations may mean that banks' long-standing plans for global consolidation of technology platforms are no longer viable, and they would need to rethink their data and technology architectures.


The extent to which a bank is subject to these effects depends on its business mix, technology strategy and the countries in which it operates. Wholesale banks that participate in a relatively small number of markets — with fewer customers and relatively flexible technology architectures — appear less concerned about data location regulations. They consider the issue to be just another type of complexity for them to manage. By contrast, retail financial institutions that participate in a large number of markets — with millions of customers and, often, monolithic technology architectures — are more affected.

POTENTIAL ACTIONS FOR FINANCIAL INSTITUTIONS

There is a range of potential actions financial institutions are considering that can reconcile burgeoning data location regulations with their business aspirations, and help them make effective use of global digital opportunities (see Figure 4):

- **Devote the resources and expertise required to create transparency and insight** into data location regulations, across regulatory agencies and markets. For example, one bank created a database of more than 1,000 data location and privacy regulations for the markets in which it participated. In some cases, when banks have achieved detailed insight into local regulations, it can help reduce their economic impact. Over time, financial institutions may be able to create shared utilities to maintain common regulatory databases.
- **Incorporate impact of data location regulations into country-level strategies.** One financial institution is explicitly considering the cost of required local technology in regional business strategies, including market exits.
- **Accelerate efforts to develop more flexible technology platforms.** This will help banks reduce the economic impact of data location regulations. One bank increased the modularity of its application portfolio, which reduced the amount of supporting technology required to maintain data in each country. Another is considering using its private cloud infrastructure to shift workloads around more efficiently — this too will lower costs in countries with

Figure 3: Correlation of Banked Population to Bank Density, 2011



Data sources: Reserve Bank of India, Organisation for Economic Co-operation and Development, World DataBank, Wikipedia, Brazil Central Bank, Bank Indonesia, World Banking Intelligence, India Central Bank, FSS Korea.

Note: 1. Population over 15 years of age, with an account in a formal financial institution.

new data location requirements. Better technology can also offer the ancillary benefit of providing deeper insights into the bank’s customer base: for example, age breakdown of customers with over US\$1 million in assets.

- **Consider working together as a sector to persuade regulators** that global operating models in financial services can benefit consumers by increasing the number of institutions competing for their business in each market. As part of this collective effort, the industry should make regulators comfortable that their national objectives will not be endangered by the presence of global models.

POTENTIAL ACTIONS FOR REGULATORS

Policy makers should continue to seek the right balance between national policy concerns and enforcement of regulation that would reduce gains of scale of global institutions. They could consider the following steps:

- **Assess the economic cost of these regulations** by working with global firms (even beyond banking) to understand the full cost of accommodating the applicable location regulations.
- **Calibrate their application carefully** by identifying firms that are drastically affected by these regulations and developing alternatives that would achieve the same goal (for example, increased security, job creation).

Figure 4: Approach to Design a Data Location Strategy



Source: Authors.

- **Ensure consistency of application of regulation** by interacting with industry bodies to drive alignment between jurisdictions, investing in building capabilities of companies related to these regulations, and managing potential conflicts in regulation across regions.

CONCLUSION

Clearly, the impact of data location regulations on financial services is significant in terms of operational complexity and cost. But data location regulation policies are legitimate attempts to address valid national policy concerns. The right balance can be achieved only if policy makers understand the economic cost of these types of regulations, calibrate their application carefully and ensure consistency of their application.

Financial institutions will need to consider a broad range of actions to reconcile regulatory compliance with their aspirations to globalize operating models, deliver innovative products and continue to drive attractive economics. Beyond the requirements of local data regulation, financial institutions also need to demonstrate their own commitment to customer data security and privacy, which could turn out to be more restrictive than that of the countries they operate in. Demonstrating such commitment will strengthen their position in dealing with national authorities, as will appropriately balancing their business aspirations with a focus on consumer impact and ethical standards.

ABOUT THE AUTHORS

James M. Kaplan is a partner at McKinsey & Company in New York. He convenes McKinsey's global practices in IT infrastructure and cyber security. He has assisted leading institutions in implementing cyber-security strategies, conducting cyberwar games, optimizing enterprise infrastructure environments and exploiting cloud technologies. James led McKinsey's collaboration with the World Economic Forum on "Risk & Responsibility in a Hyper-Connected World," which was presented at the forum's recent annual meeting in Davos. He has published on a variety of technology topics in the *McKinsey Quarterly*, the *Financial Times*, *The Wall Street Journal* and the *Harvard Business Review* Blog Network.

Kayvaun Rowshankish is a partner at McKinsey & Company in New York, where he is a member of the Business Technology Office. He serves financial institutions (universal banks, asset managers, investment banks and securities firms) on a broad range of operations and technology related issues. Kayvaun is a global leader of the Data & Analytics practice in Banking. Prior to joining McKinsey, Kayvaun spent 11 years at Diamond Management and Technology Consultants (acquired by PwC), where he was a partner in their capital markets practice. Kayvaun graduated from Surrey University with a B.A. in business studies and he received an M.B.A. from Cass Business School.

CHAPTER THREE: INTERNET OPENNESS AND FRAGMENTATION: TOWARD MEASURING THE ECONOMIC EFFECTS

Sarah Box

Copyright © 2016 by the Organisation for Economic Co-operation and Development

ACRONYMS

BSA	The Software Alliance
ccTLDs	country-code top-level domains
CDNs	content distribution networks
gTLD	generic top-level domain
GVCs	global value chains
ICT	information and communication technology
IP	Internet Protocol
ISP	Internet service provider
IXP	Internet exchange point
Mbps	megabits per second
OECD	Organisation for Economic Co-operation and Development
R&D	research and development
SMEs	small and medium-sized enterprises
TLDs	top-level domains

INTRODUCTION

Internet openness and Internet fragmentation are often portrayed as opposing forces struggling for ascendancy. If Internet openness wins, we have a world of global connections and freedoms. If Internet fragmentation wins, we have a world of silos and closed doors. This kind of scenario implies significant economic consequences and people understandably want to know what exactly, and how large, these consequences might be.

Through its work on the economic and social benefits of Internet openness, the Organisation for Economic Co-operation and Development (OECD) is attempting to bring new evidence to the debate. This is important because, in reality, the issue is not black and white: openness is not indisputably good and fragmentation is not indisputably bad. Governments need more nuanced information to allow policy choices that optimize the benefits of Internet openness while addressing valid concerns for digital security and privacy. Progress must be made in understanding the strength and direction of the relationship between Internet openness and governments' ultimate economic goals — such as enhanced trade, innovation and entrepreneurship — and how Internet openness itself is affected by policy and private sector actions.

Analysis of Internet openness quickly meets a practical stumbling block: how do we measure it — or, indeed, measure Internet fragmentation — when the concept of Internet openness itself is so broad, encompassing technical, economic, political and societal aspects? To make headway, the OECD chose to focus efforts on better understanding and measuring global data flows on the

Internet, as an initial indicator of Internet openness. From this starting point, it has begun building a picture of global data flows and laying out a path for future analyses (OECD, forthcoming 2016). This chapter presents an excerpt of that work.¹ It describes the benefits of Internet openness for international trade, innovation and entrepreneurship, and presents initial steps to better measure the global data flows enabled by Internet openness.²

INTERNET OPENNESS AND INTERNATIONAL TRADE

There is a growing literature on the positive effects of the Internet on trade and the potential costs of policies (notably on data localization) that introduce frictions to “business as usual” data flows on the Internet. Internet openness facilitates international trade for existing businesses by making it easier for the supplier to connect with existing consumers who are located beyond the borders of the supplier’s home country (or countries) and by improving logistics control. Openness can also boost trade by providing access to a wider customer base via e-commerce. And it enables new firms to enter more geographic markets and, for the most efficient ones, to enter global value chains (GVCs). At the same time, Internet openness and digitization make it possible to complete transactions and deliver products, services and payments faster and more efficiently by replacing some physical trade with online trade — for example, in books and music, or with more complex products via online shipment of designs followed by local production.

GVCs are central to the trade and Internet story. Behind aggregate trade data lie a huge number of intermediate trade flows, with inputs sourced globally and stages of production shifting from location to location to complete a final product. Both goods and services may be produced in GVCs — electronics and cars are common examples where design, raw material, production and marketing inputs are spread across countries, but aircraft, clothing, film animation, law briefs and medical advice are also created in GVCs. The rise of GVCs has been made possible in part by technological advances, notably the information management systems that allow firms to coordinate their participation in GVCs. The combination of GVCs and the Internet has not only enabled firms in developing countries to more easily engage in international trade (by specializing

1 The OECD’s work on Internet openness is being undertaken in the context of the 2016 OECD Ministerial Meeting on the Digital Economy: Innovation, Growth and Social Prosperity, to be held in Mexico in June. See www.oecd.org/sti/dep-ministerial-2016.htm. The Ministerial has four central themes: Internet Openness and Innovation; Building Global Connectivity; Trust in the Digital Economy; and Jobs and Skills in the Digital Economy.

2 This chapter should be read in conjunction with the chapter *A Framework for Understanding Internet Openness* by Jeremy West.

in one stage of a chain, such as auto electronics), but also through the use of digital platforms provided by small and medium-sized enterprises (SMEs) to enable even tiny firms (micro-multinationals³) to connect with global suppliers and purchasers.

Seamlessly moving potentially large amounts of data across countries is an essential part of supporting intermediate and final trade flows and allowing firms to participate in GVCs. In other words, given the pervasiveness of GVCs, reductions in Internet openness could create significant impediments to trade. Small frictions may multiply into large barriers, especially if production is split into stages that entail numerous border crossings where imposed frictions multiply. The Swedish National Board of Trade (2015, 14-15) suggests that policies such as data localization requirements could lead a firm to reorganize its GVC, either moving or closing parts of its operations, with service to end-users being restricted in some cases. Stephen Ezell, Robert D. Atkinson and Michelle Wein (2013, 46-47) make a similar point, noting that localization barriers to trade, including restrictions on data, undermine firms' ability to participate in global networks because the barriers raise costs and reduce technology diffusion. The Software Alliance, more commonly known as the BSA, additionally highlights the trade-dampening effect of country-specific technology standards and other forms of "digital protectionism," such as nationally oriented information technology procurement (BSA 2014).

Internet openness is especially important for enabling smaller firms to engage in international trade. Jessica R. Nicholson and Ryan Noonan (2014, 8) comment that while localization requirements can make cross-border trade difficult for large companies, they may make it "practically impossible for small businesses that cannot afford to implement separate systems and standards in every country in which they do business." Moreover, these firm-level impacts can sum to significant negative outcomes for countries. James M. Kaplan and Kayvan Rowshankish (2015) note that as banks reduce their operations in countries with more stringent data regulations, financial services will grow more slowly, with potentially adverse consequences for development. There are also more general concerns that policies enacted to reduce Internet openness could create a "slippery slope" for additional interventions and possibly non-tariff barriers, such as local content requirements or efforts to promote "indigenous innovation" via intellectual property right restrictions.

3 The term micro-multinational is not well defined and should not be automatically equated with small multinational enterprises. Micro-multinationals may simply be small exporters, whereas multinational enterprises typically comprise "companies or other entities established in more than one country and so linked that they may co-ordinate their operations in various ways" (OECD 2008, 12). Ann Mettler and Anthony D. Williams (2011) discuss micro-multinationals in terms of start-ups, typically small, service-driven companies.

Ezell, Atkinson and Wein (2013, 38) see a risk that the contravention of the rules and spirit of the global trading system would lead to a decay where "every country is incentivized to cheat, the competition becomes cutthroat, and the global economy suffers."

INTERNET OPENNESS, INNOVATION AND ENTREPRENEURSHIP

The Internet, as a connector on a massive scale, provides the opportunity to share, access and coordinate knowledge in ways previously not possible. Knowledge sharing was the impetus behind the creation of the Internet, albeit among an initially small group of research institutions, and research-oriented knowledge-sharing networks running on the Internet remain. These help facilitate collaborative research on a global scale, with publications, patents, researchers, and academic and research institutions taking on international dimensions and drawing benefits from cross-border knowledge flows. Firms, too, leverage the Internet to share knowledge, from multinationals with diverse research and development (R&D) and production locations to small firms tapping into local universities and research institutions. And the general expansion of access to knowledge (for example, via Google searches, Wikipedia, YouTube or online education sites) to a broader range of people can also stimulate innovation. Joshua Meltzer (2015, 92) states:

The Internet has provided an opportunity for people to connect and share ideas in a space and time essentially free of transaction costs. Significantly, it has been the open nature of the Internet — the freedom to connect, share information and exchange ideas — that has underpinned the innovation which has created new businesses such as those based on social networking and crowd funding.

The Internet also provides a platform for innovation, open to anyone who wishes to leverage it for their own venture. Several aspects of this are frequently mentioned — first, that the Internet enables "innovation at the edges"; second, that it enables "serendipitous" (or unexpected) innovation; and third, that it allows "permissionless" innovation. The term "innovation at the edges" references the Internet's end-to-end design principle, whereby the core network provides general purpose system services (sending packets of data) and is indifferent to the various applications that may be implemented in software on computers attached to the "edge" of the Internet (Blumenthal and Clark 2001). This end-to-end feature makes the Internet flexible, general and open to innovative new applications. These innovations can challenge the status quo and can bubble up from unexpected quarters (hence the idea of serendipity), including from very small firms. Finally, permissionless

innovation captures the idea that market entrants need not seek approval prior to launching lawful new services, and that this lack of gatekeeping leads to a flourishing market for ideas, be it through social networks or through promoting innovation around new devices and services. Leslie Daigle (2015, 9) points to the creative destruction built into the Internet, saying, “Systemically, the Internet supports and fosters approaches that are useful; old, outdated or otherwise outmoded technologies die away.”

As a source of inputs to entrepreneurs and established firms, the Internet is also becoming increasingly valuable, offering a conduit to finance, services and marketplaces. In a way, the Internet is taking outsourcing to its extreme, allowing firms to fully concentrate on their competitive advantage. This not only benefits existing firms by improving efficiency and providing headspace for new innovative activities, but also makes it easier for entrepreneurs to muster the resources to take their ideas through to commercialization. The new phenomenon of micro-multinationals, for instance, is underpinned by the availability of business services via Internet platforms (Mettler and Williams 2011), and SMEs can also reap significant rewards from boosting their digital savvy (Mettler and Williams 2012). Firms can design, develop and deliver their products and services worldwide thanks to Internet-based crowd financing, digital utilities, professional services, micro-manufacturing, innovation marketplaces and e-commerce platforms.

Lastly, the information and communication technology (ICT) sector itself is a generator of innovation, offering increased computer power and performance and new tools. This sector forms part of the economic constellation around the Internet and both nourishes and feeds off the economic and social activity enabled by the Internet. The ICT sector was relatively resilient to the 2007–2009 global economic crisis, although it has yet to regain its pre-crisis levels in some countries, and is an important venue for R&D and patenting. Advances in ICT will underpin data-driven innovation — for instance, the main enablers of the Internet of Things are big data, the cloud, machine-to-machine communication and sensors (OECD 2015a, 244).

But all this relies crucially on Internet openness — free flows of data and information, accompanied by trust in the network, are essential for the Internet to contribute to innovation and entrepreneurship. In a recent study, young entrepreneurs in Group of Twenty countries identified international mobility of data accompanied by adequate protection of personal data as a key issue, saying that this was “one of the success factors of entrepreneurs who develop international businesses, and a critical element for entrepreneurs to get access to the right data” (Accenture 2013, 36). Commentators have argued that innovation in industries such as ICT, energy, life sciences, aerospace and scientific instruments could be especially impeded by limits to data mobility, since such industries do best

serving large markets in a competitive environment (Ezell, Atkinson and Wein 2013). Limiting scale economies enables weaker firms to remain in the market, thus reducing returns to more efficient firms and eroding their ability to invest in innovation. At the same time, security and privacy standards are necessary to support innovation on the Internet; for example, in Estonia, the X-Road data exchange framework enables access to publicly held data in a high-trust environment and has spawned the development of numerous new Internet businesses, including Skype (Hofheinz and Mandel 2014).

MEASURING INTERNET OPENNESS

Specific studies on Internet openness are still scarce and there is much scope for improving quantitative evidence on the links between Internet openness and economic indicators such as trade and innovation. But, as noted earlier, the concept of Internet openness is so broad that measurement is a significant challenge.

Existing studies of the Internet’s macroeconomic impact have typically used various proxies of Internet presence, including adoption indicators (such as broadband penetration rates), economic indicators (such as network investment) and technical indicators (such as Internet Protocol [IP] addresses per capita). Each of these proxies has limitations, one being lack of insight into how people, firms, industries or regions actually make use of the Internet (OECD 2012). Unfortunately, these proxies are also imperfect measures of Internet openness, as they essentially focus on access and availability.

Quantitative studies of the Internet from a digital trade angle have typically used proxies of data flows for their analysis. On the face of it, using data flow information as a measure of Internet openness has merit. If the essence of the Internet is to facilitate movement of data/information/knowledge, for whatever purpose, then measuring flows of data could shed light on current levels of openness, even if the economic value of the data flows is unknown. Changes in flows could then be related to changes in trade and other variables on the one hand, and changes in policy or other factors on the other hand (assuming we could construct robust policy indicators). In addition, as many of the risks to Internet openness are occurring at the level of data flows, measuring this aspect would be highly relevant.

However, the proxies of data flows used to date also have drawbacks:

- As Paul Hofheinz and Michael Mandel (2015) point out, using official statistics (such as trade data related to digital activity) essentially underestimates the size of cross-border data flows, because not all flows are monetized.

- While looking at the bits and bytes themselves is another option, information on the capacity of the infrastructure (such as TeleGeography statistics [McKinsey Global Institute 2014]) does not inform us of actual data flows.
- Adding capacity usage estimates or traffic estimates can bring us closer to actual data flows, but such estimates (for instance, Cisco global IP traffic forecasts [Hofheinz and Mandel 2014; 2015]⁴) do not differentiate where the traffic is coming from or going to — i.e., whether start and end points are local or cross-border — or the type of flows.

In one of the few studies that have approached Internet openness more directly, Dalberg (2014) chose to use Freedom House’s Freedom on the Net index to look at the economic benefits of Internet openness. This index is based on qualitative assessments and surveys, and measures the level of Internet and digital media freedom in three areas: obstacles to access (such as regulatory obstacles for Internet service providers [ISPs]); limits on content (for example, instances of filtering); and violations of user rights (such as state surveillance). However, Dalberg considered that the limited time series and country coverage did not allow statistically significant causal relationships to be established; indeed, one of its key conclusions was to urge stakeholders “to establish standard and universally measurable indicators of Internet openness” (ibid., 50).

Other efforts are emerging along the lines of the Freedom on the Net index that group together various indicators of Internet activity, including aspects that touch on Internet openness. For instance, the Boston Consulting Group’s e-Friction Index agglomerates 55 indicators to indicate the ease with which people can participate in the Internet economy (Zwillenberg, Field and Dean 2014). The e-Friction Index could perhaps be interpreted as an openness index, although some of the indicators (such as company-level technology absorption or financing through local equity market) are relatively upstream from practical Internet openness; furthermore, there are significant data gaps. Another effort to draw together a variety of indicators on Internet trends comes from the Berkman Center for Internet & Society, whose Internet Monitor research project aims to shed light on Internet content controls and Internet activity worldwide.⁵ As well

4 Hofheinz and Mandel’s (2015) concept of “digital density” (the amount of data used per capita in an economy) as a proxy of data usage is based on Cisco IP traffic forecasts for major countries, which are built on a series of estimates of user numbers, adoption rates, minutes of usage and bitrates to obtain a per-month traffic estimate (Cisco 2015a; 2015b). Hofheinz and Mandel (2015) acknowledge that using this as a proxy for consumption of cross-border data flows is a leap, but propose this measure gets closer to data usage than other measures of cross-border data flows.

5 See <https://thenetmonitor.org/> for further details on the Berkman Center’s initiative.

as an index related to Internet access and infrastructure, a “dashboard” was recently launched that incorporates data on traffic, cyber attacks and website availability, among other indicators.

However, it remains the fact that there is no easy off-the-shelf solution to measuring Internet openness. As such, one goal of the OECD’s work is to push the data boundaries by collecting and using data obtained from companies with global reach to provide a new perspective on global data flows across the Internet. Eventually, this work should facilitate analyzing the effects of Internet openness at a more general level than is found in case studies of individual firms or situations, and thus should help reinforce the evidence base available to policy makers.

At the time of writing, the OECD had analyzed aggregate information related to Google searches and YouTube views (see Box 1). Google and YouTube usage provide insight into the website domains that users in a country visit via Internet search, and where YouTube content is watched. While the information does not give a sense of volumes (as it was expressed in percentages), some 240 countries are covered in the tables the OECD analyzed, enabling the exploration of interlinkages. At this stage the analysis has mainly focused on OECD countries plus its key partners and accession countries (Brazil, the People’s Republic of China, Colombia, Costa Rica, India, Indonesia, Latvia, Lithuania, Russia and South Africa).

Key findings and lessons from the information analyzed are highlighted below. In interpreting the results, it is important to bear in mind the following factors:

- A ccTLD for a website does not necessarily imply that the content is hosted within that country. For instance, you do not need to be based in New Zealand to register a .nz domain name, and the domain name is not required to be hosted in New Zealand.⁶
- Indeed, some ccTLDs have no substantive linkage to the country at all and instead are used much like a generic top-level domain (gTLD). Examples include Belize (.bz), the Cocos (Keeling) Islands (.cc), the Federated States of Micronesia (.fm), Lao People’s Democratic Republic (.la), Montenegro (.me), Niue (.nu), Samoa (.ws), Sint Maarten (.sx), Tokelau (.tk), Tonga (.to) and Tuvalu (.tv).
- A gTLD for a website cannot be matched to a particular country, either in terms of “owner” of the site or where the content is hosted, as these domains are available

6 See more information at the .nz Domain Name Commission at <https://dnc.org.nz/the-commission/faq>.

Box 1: Google Data Specifications

The OECD analyzed four tables of information, related to Google searches and YouTube watch time, as follows:

Source 1: Google Search — Focus on User Country

A table of 240 countries¹ (including 1 “zz” category where the country of the user could not be determined) by 101 top-level domains (TLDs — comprising 87 country domains, 13 generic domains and 1 “other” category), showing the percentage of clicks on search results by users of a particular country searching on Google (all domains) that landed on websites of each TLD.² This allows us to see, for instance, that in 2014, five TLDs (.com, .au, .org, .net and .uk) accounted for 96.11 percent of Australian users’ Google search result clicks, with the remaining 3.89 percent of clicks going to a variety of landing page TLDs. User locations were based on IP addresses.

Time span: 2007–2014 (eight years) for most countries in the table.

Source 2: Google Search — Focus on Landing Page TLD

A table of 240 countries (including 1 “zz” category where the country of the user could not be determined) by the same 101 TLDs, showing the percentage of clicks on search results related to each landing page TLD that come from users of a particular country who are searching on Google (all domains).³ This allows us to see, for instance, that in 2014, 25.35 percent of clicks received by .com landing page domains via Google search results came from users in the United States. User locations were based on IP addresses.

Time span: 2007–2014 (eight years) for most countries in the table.

Source 3: YouTube — Focus on Country of Uploader

A table of 240 uploading countries by 240 watching countries, allocating the percentage share of watch hours of an uploading country’s YouTube videos across each watching country. There is additionally a “zz” category where the countries of uploading user and watcher could not be determined.⁴ This allows us to see, for instance, that in 2014, 18.23 percent of the watch hours for videos uploaded by users from Spain were by users located in Mexico — the second-highest watch hour share after Spanish viewers (at 23.44 percent). The locations of uploading users were user-specified, and those of watching users were based on IP addresses.

Time span: 2010–2014 (five years) for most countries in the table.

Source 4: YouTube — Focus on Watching Country

A table of 241 watching countries by 250 uploading countries (each including a “zz” category where the countries of uploading user and watcher could not be determined), allocating the percentage share of a country’s YouTube watch hours across different YouTube video uploading countries.⁵ This allows us to see, for instance, that in 2014, Slovenian users spent 1.61 percent of their YouTube watch hours on videos uploaded by users in Italy. The locations of uploading users were user-specified, and those of watching users were based on IP addresses.

Time span: 2010–2014 (five years) for most countries in the table.

1 References to “country” should be read to include all geographic areas with two-digit country code top-level domains (ccTLDs) in the tables. These include the 193 member states of the United Nations as well as other territories.

2 As the information is in percentages, it is not possible to say how large the “zz” user category is compared to other user countries. However, the share of user clicks going to the “other” category are typically small; for instance, for all OECD key partner and accession countries, except for Luxembourg, the shares of user clicks going to the “other” category are less than one percent. In Luxembourg’s case, 13–17 percent of clicks went to “other” over the 2007–2014 sample period.

3 In this table, it is not possible to say how large the “other” category is compared to the other TLDs, but we can see that the “zz” user category makes up less than one percent of clicks on TLDs in the majority (84 percent) of cases. Over the eight-year period, .co (Colombia), .id (Indonesia), .in (India), .ir (Islamic Republic of Iran), .pk (Pakistan), .sa (Saudi Arabia) and “other” saw the most frequent incidences of a high “zz” user share.

4 In this table, it is not possible to say how large the “zz” category is as an uploading country, but we can see that “zz” as a watcher accounts for less than one percent of watch hours for any country’s YouTube videos in the majority (94 percent) of cases, with this share typically decreasing over the sample period. The most frequent incidences of a high “zz” watcher share were for .al (Albania), .ir (Iran), .mc (Monaco) and .mk (Former Yugoslav Republic of Macedonia).

5 In this table, it is not possible to say how large the “zz” category is as a watcher country, but we can see that “zz” as an uploader has accounted for a steadily decreasing share of each country’s watch hours over the sample period. In 2010, the share of watch hours going to “zz” YouTube videos reached 15 percent in some cases (Iran and Japan), but by 2014, the share was below or close to one percent in all cases.

for registration by Internet users worldwide (albeit with some restrictions for some domains⁷).

- The network architecture of the Internet, the extensive use of data centres (“the cloud”) and the growing presence of content distribution networks (CDNs) mean that the physical route taken by data may bear little resemblance to a straightforward bilateral flow between two countries.

INSIGHTS FROM INFORMATION ON GOOGLE SEARCHES

The Google search information from source 1 in Box 1 shows that Internet users differ widely in the extent to which they select results in their own country’s domain. For instance, in 2014, 67 percent of Google search clicks by users in Poland led to .pl domains, whereas only 13 percent of search clicks by users in Korea led to .kr domains (see Figure 1). The United States is an exceptional case; for historical reasons, gTLDs such as .com were preferred to the .us domain, which was commercially marketed at a later stage, and just 0.66 percent of US users’ Google search clicks went to .us websites in 2014.

Accompanying this diversity is an almost uniform trend of users increasingly accessing content outside their countries. With the exception of Canada, Estonia, France, India, Ireland and Sweden, all countries experienced a decline in the share of Google search clicks going to their own ccTLD between 2007 and 2014. These findings might suggest a geographically wider variety of content being accessed, increased cross-country information and knowledge exchange, and potentially an increase in actual cross-border data flows, subject to the caveats mentioned earlier.

The extent to which these patterns are accompanied by changes to Google’s search algorithms is an interesting question. Google’s algorithms rely on over 200 “signals” to help guess what the user might be looking for in their search, including terms on websites, content freshness, the user’s region and PageRank (a measure of how authoritative a webpage is).⁸ An increased internationalization of the content accessed by Google users could reflect many factors and developments. It is possible that the queries issued by users over time relate to more international topics (i.e., a change in the “query mix”), thereby leading to more international results surfacing. Even for an unchanged query mix, it is possible that users over time become more interested in international sources, seeking them out in search results; this could potentially be accompanied by

Google’s algorithms taking account of this preference in the composition of search results. In addition, the shape of the underlying Internet is ever-changing, and to the extent that the growing number of web pages “internationalize” this base, one would expect this change to be reflected in Google’s index as well. Irrespective of the precise explanation, the fact remains that many users are increasingly looking beyond their own country content.

The information on gTLDs show that a significant share of users’ search clicks go to sites with a .com domain. In fact, in every country, .com domains were the most or second-most common result click, along with the country’s ccTLD (with the exception of China, Korea, Luxembourg and the United States, where the .com domain was accompanied by .hk, .net, “other” and .org, respectively, in the top two). Thirteen gTLDs were included in the Google search information — .com, .org, .net, .edu, .info, .gov, .biz, .cat, .mobi, .xxx, .mil, .name, .int — with .com, .net and .org uniformly the top three gTLD clicks and cumulatively accounting for over 50 percent of search result clicks in 27 of the countries in 2014 (see Figure 2).

The importance of language/culture and geographic proximity can be observed in the search information. Proximate countries and those with a common language are typically among the top 10 ccTLDs in a country’s search result clicks. For example, Chilean users click on results in the Spanish, Argentinian, Mexican, Colombian, Peruvian and American ccTLD spaces, while Swiss users click on results in the German, French, Italian, UK and Austrian ccTLD spaces. This behaviour is consistent with international trade models for goods and services that show that “gravity” — as measured by proximity, common language and so on — is an important factor driving trade links, although there may also be other effects in operation.


At the same time, the usage of the generic ccTLDs is also notable. While Tonga and Tuvalu might seem logical search result clicks for users in Australia and New Zealand — Pacific neighbours and home to immigrant communities — it is less obviously the case for Estonia and Israel, and the widespread appearance of these generic ccTLDs in top 10 search result click lists underscores the lack of a one-to-one relationship between ccTLDs and their “home countries.” For instance, Tuvalu’s ccTLD is often used by media companies (the .tv domain name having clear marketing value). Nevertheless, the share of total search result clicks received by such TLDs is typically small since, as clearly illustrated in Figure 2, gTLDs account for a significant share of total user clicks.

The Google search information from source 2 in Box 1 suggests that most website ccTLDs have a highly concentrated user base, accompanied by a long tail of user countries, each with tiny shares of total search result clicks. Taking the full sample of ccTLDs included in the table (excluding those that are clearly used in practice as

⁷ See the Internet Corporation for Assigned Names and Numbers’ list of TLDs and registrars at www.icann.org/registrar-reports/accredited-list.html.


⁸ See www.google.com/insidesearch/howsearchworks/algorithms.html.

Figure 1: Share of Google Search Result Clicks Leading to Sites with Own ccTLD (2007 and 2014)



Source: OECD calculations, based on information from source 1 (see Box 1).
 Note: Data on Luxembourg (.lu) as a search domain was available in the table.

Figure 2: Share of .com, .org and .net in Search Result Clicks by Country (2014)



Source: OECD calculations, based on information from source 1 (see Box 1).

gTLDs), 41 of 75 ccTLDs received 95 percent of search result clicks from four or fewer user countries in 2014. These were typically the country of the ccTLD plus proximate countries (either geographically or via cultural/language similarities). For instance, users from Israel and the United States accounted for over 95 percent of search result clicks to websites with Israel's ccTLD (.il), while users from South Africa, the United States and the Netherlands accounted for over 95 percent of search result clicks to websites with South Africa's ccTLD (.za). Most OECD countries received 95 percent of search result clicks from six or fewer user countries.

However, some ccTLDs have lower levels of concentration, although still with the long tail. OECD countries that stand out in this respect include Spain (12 user countries accounted for 95 percent of search result clicks in 2014), as well as Sweden, the United Kingdom, the United States and Iceland (20, 21, 27 and 50 user countries, respectively). Mexico and Colombia accounted for a significant share of Google search result clicks to websites with Spain's ccTLD (.es), followed by a number of other South American countries, plus the United States, Germany and India. The wide range of user countries behind search result clicks to websites with the United Kingdom ccTLD (.uk) is perhaps reflective of the United Kingdom's historic Commonwealth links as well as its status as a global hub.

The user base of gTLDs is unsurprisingly less concentrated than that of ccTLDs, matching their greater global availability. But one interesting observation is the variety of user countries for the gTLD .edu, which is available only to US post-secondary institutions that are accredited by an agency on the US Department of Education's list of Nationally Recognized Accrediting Agencies.⁹ The Google search result clicks could be interpreted as mirroring the international attractiveness of the United States as an education destination. Users from the United States accounted for almost 71 percent of search result clicks to .edu domains in 2014; users from 27 other countries (shown in Figure 3) then accounted for a further 24 percent of the clicks.

INSIGHTS FROM INFORMATION ON YOUTUBE WATCH HOURS

YouTube is a platform for user-generated video content, from music to do-it-yourself bicycle repairs, from professional to amateur. It has been credited as a source of ideas and cross-fertilization.¹⁰ The YouTube information in sources 3 and 4 (see Box 1) do not distinguish between

types of content, but they do provide an aggregated picture of the viewing patterns of YouTube users.

Figure 4 shows a wide variation in the extent to which content is viewed outside the country in which it is uploaded. In 2014, for instance, 85 percent of the watch hours for videos uploaded by users in Japan were from users located in Japan. Toward the other end of the scale, just eight percent of the watch hours for videos uploaded by users in Australia and Canada were from users located in those countries. For both Australia and Canada, users in the United States accounted for the largest share of watch hours for Australian- and Canadian-uploaded content (27 and 37 percent, respectively). US users were the second-largest share of viewers of Japanese YouTube content, with almost three percent of watch hours.

Figure 4 also shows how, for more than half of the examined countries, dispersion of content is becoming increasingly international. In the United States, for example, the share of watch hours for US-uploaded content accounted for by US users fell from 42 percent to 35 percent over the period 2010–2014. After US users, the top watchers of US-uploaded YouTube content in 2014 were the United Kingdom, Vietnam, Mexico, Canada, Russia, Japan, Australia, Brazil, Germany and Turkey, in that order. In contrast, Japan, Brazil, Turkey and others saw an increase in the share of local watchers in watch time for their content between 2010 and 2014. In some instances, this may be because the amount of local content being produced is increasing and attracting new local users; this, in turn, may be related to the penetration of smart phones, which offer another way to capture and view content.


The range of countries among watchers of a country's content sometimes points to the importance of a common language. For instance, YouTube content uploaded by Spanish users in 2014 obtained its highest share of watch hours from local viewers (23.4 percent), followed by Mexico (18.2 percent), Argentina (9.1 percent), the United States (6.1 percent), Chile (5.6 percent), Colombia (5.3 percent), Peru (3.6 percent), Venezuela (2.5 percent) and Ecuador (2.4 percent). Proximity and historical links can also be observed — in France, for instance, the highest share of watch hours of content uploaded by French users in 2014 came from France (50.5 percent), followed by the United States (5.1 percent), Belgium (4.3 percent), Canada (3.0 percent), Morocco (2.6 percent) and Algeria (2.0 percent).

Focusing on what people watch, source 4 shows that for the most part, the share of any country's watch hours spent on another country's YouTube content is numerically small (i.e., less than one percent), implying that in aggregate, people are taking a smorgasbord approach — a little bit of lots of things. However, there are three instances where this is not the case:

⁹ The .edu domain's sole registrar is Educause, an association for information technology in higher education. Eligibility for the .edu domain name is restricted. See <http://net.educause.edu/edudomain/eligibility.asp>.

¹⁰ See, for example, McNeil (2013).


Figure 3: Top Users of .edu gTLD, Measured by Share of Search Result Clicks (2014)



Source: OECD calculations, based on information from source 2 (see Box 1).

Note: The United States is excluded from this figure. US users plus the 27 countries in the chart accounted for 95 percent of Google search result clicks to websites with a .edu TLD.

Figure 4: Views of YouTube Content Uploaded by Users in Own Country (% of Total Watch Hours for Country's Uploaded Content)



Source: OECD calculations, based on information from source 3 (see Box 1).

- All watching countries spent 10 percent or more of their watch hours on US-uploaded content, with 20 countries spending more than 50 percent of their watch hours on US-uploaded content (aside from the United States itself, these were Caribbean island nations plus Antarctica,¹¹ Bermuda, the Marshall Islands and several US island territories).
- Some countries' consumption of local content accounts for very high shares (over 50 percent) of total watch hours. Brazil stands out as a large consumer of its own content — 72 percent of its watch hours are on Brazil-uploaded content. Indian users also spend more than half their watch hours on local content (almost 58 percent). Other countries in this category are Japan (65 percent), Korea (62 percent), Poland (55 percent) and Thailand (66 percent).
- Certain countries' content more regularly accounts for a high share of watch hours in other countries. Spain, France and the United Kingdom stand out, with their content accounting for 10–50 percent of a relatively large number of watching countries' total watch hours (20, 38 and 45 countries, respectively). There are clear language and historical links — for instance, the countries for which French content accounts for 10–50 percent of watch hours are Algeria, Belgium, Burkina Faso, Burundi, Benin, Cameroon, Central African Republic, Chad, Comoros, Congo, Côte d'Ivoire, Democratic Republic of the Congo, Djibouti, French Guiana, French Polynesia, Gabon, Guadeloupe, Guinea, Haiti, Luxembourg, Madagascar, Mali, Martinique, Mauritania, Mauritius, Mayotte, Monaco, Morocco, New Caledonia, Niger, Réunion, Saint Pierre and Miquelon, Senegal, Switzerland, Togo, Tunisia, and Wallis and Futuna (as well as France itself).

Aside from these patterns, there are also some individual cases that stand out. For instance, Indian content accounts for more than 10 percent of watch hours in several Middle Eastern countries (for example, 24 percent of watch hours for the United Arab Emirates, 15 percent for Bahrain, 12 percent for Kuwait, 15 percent for Oman and 22 percent for Qatar). Fijian users also spend a significant share of watch hours on Indian content (26 percent). This may be due to past and recent immigration patterns that have created significant Indian communities in these countries

¹¹ The Antarctica (.aq) TLD is administered by the Antarctica Network Information Centre Limited located in New Zealand. The .aq domain name is available to government organizations who are signatories to the Antarctic Treaty and to other registrants who have a physical presence in Antarctica. Due to the special nature of the Antarctic environment, the registrar considers a “physical presence” to include unattended installations owned or operated by the registrant and short-term visits to the ice by the registrant or its employees. Enthusiastic consumption of US-uploaded YouTube content may be partly due to the large US base in Antarctica (McMurdo Station).

and/or to the creation of content in India that particularly appeals to Middle Eastern users.


Given the factors above, a country's watch hours typically display a long-tailed pattern, much like that of the earlier information on Google searches, where most watch hours are dedicated to content from a small group of countries and the remainder of watch hours are accounted for by small amounts of many countries' content. Four country examples are presented in Figure 5. In each case, the pie chart specifies the uploading countries (in descending order of importance) that together account for 80 percent of watch hours, with the remainder of sources aggregated as “other.” It shows that for Italy, 10 countries accounted for around 80 percent of Italian YouTube watch hours in 2014, although within that a large chunk was local Italian and US content. South Africa also had 10 countries accounting for around 80 percent of its watch hours, in this case led by the United States, the United Kingdom and then local content. Eight countries accounted for 80 percent of Colombians' watch hours in 2014; this time was more evenly spread among US, Spanish, local and Mexican content. Turkey stands out, with just four countries accounting for 80 percent of watch hours, namely Turkey, the United States, Germany and the United Kingdom.

The table from source 4 also provides the possibility to observe how watch patterns have changed over the period 2010–2014 for individual countries. As an example of this type of analysis, Box 2 looks at seven African countries — Cameroon, Ghana, Kenya, Malawi, Nigeria, Rwanda and Tanzania. Africa was the last continent to achieve Internet connection and is still in the relatively early stages of expanding access and coverage to its population. It is interesting to see that all countries in this sample have experienced an increase in the share of watch hours attributed to locally and proximately uploaded YouTube content, although the absolute shares differ widely, doubtless reflecting their different stages of digital development.

The international sharing of YouTube content is clearly a facet of global knowledge and information flows, but its value is likely to depend greatly on the content in question, as well as on how economic and social value is measured. Subject to data availability, future work could usefully explore different categories of content, distinguishing, say, education content from other content.

FROM DESCRIPTION TO MEASUREMENT?

Because of the geographic fuzziness of the information sources analyzed here, using them as a stand-alone proxy of global data flows and linking them to data on trade, innovation and other economic indicators would be misleading. In particular, the fact that gTLDs cannot be given a geographic tag makes the use of the search data to proxy data flows on a country-by-country basis unsuitable.

Figure 5: 2014 YouTube Watch Hour Patterns: Whose Content Are They Watching?

Source: OECD calculations, based on information from source 4 (see Box 1).

With .com domains representing over 40 percent of search result clicks in 2014 in 20 OECD countries (over 80 percent in the United States), for example, this loss of geographic information is significant. Added to this is the lack of one-to-one relationships between ccTLDs and the location of content. While the start and end points of data flows are clearer for the YouTube information reviewed, both it and the search information have the common problem that the actual route of data flows (and thus the interdependence of global connections) is hidden behind the bilateral data points in the tables reviewed here.

However, comparing patterns in the tables with information related to infrastructure can provide additional insights into data flows and give some pointers for the direction of work. In short, Internet infrastructure has both influenced and evolved around data flows, and continues to do so in response to market and regulatory imperatives. For instance, the growth of heavy content and consumer demands for speed and quality mean that for some types of data flows there is a clear economic case for data to stay as local as possible. One example of this might be software updates, where the same content is being downloaded multiple times and where the balance of transit costs, speed/quality outcomes and storage costs makes it sensible to shift the content close to the consumer. At the same time, there remain data flows that do not lend themselves easily to localization near the customer — they may be more unique in terms of content and need to traverse regional, if not global, networks on a constant basis. One example might be financial and logistics information flows associated with international trade.

Measures and interpretations of data flows may need to be nuanced to account for different contexts. The following discussion expands on this idea and proposes some next steps.

Location, Location, Location

The determining factor in identifying Google search destinations (and thus data flows) is where the site is hosted, and for some ccTLDs this is predominantly offshore. Figure 6 shows to what extent countries hosted the content of their ccTLD domain in 2013. It reveals that most OECD countries host at least half the content associated with their ccTLD, but there is nevertheless a wide range of outcomes, from Korea hosting almost 97 percent of .kr sites to Greece hosting just 19 percent of .gr sites. This underscores the strong global nature of the digital economy and its associated data flows. For example, 54 percent of .pt sites were hosted in Portugal in 2013. This implies that perhaps half the time, a “local” search click to a .pt website actually entailed cross-border data flows. At the same time, Portugal also hosts foreign content (in fact, in absolute terms, Portugal hosts more foreign sites than local .pt sites), thus a share of “foreign” search clicks will stay local.¹²

¹² It is possible that the data underestimates locally hosted sites, for example, in cases where content may be presented in a national and international version — say, when a newspaper hosts a site in the country for local users and has another abroad in a location close to its international readership — or where CDNs are used to distribute data. In each case, these would have shown up as hosted outside the country in the data set (OECD 2014).

Box 2: YouTube in Africa – A Peek at Watch Patterns

Source 4 (see Box 1) allows an analysis of YouTube watch patterns across a wide range of countries — too wide for this chapter to give attention to all interesting cases. However, given Africa’s status as a catch-up continent on Internet connection and usage, the table below presents some simple statistics on the change in YouTube patterns in the period 2010–2014 and current watch patterns for seven countries.

There are large differences between the countries in the share of local content watched in 2014, but all showed growth in this share from 2010 to 2014. Nigeria has the strongest local following, perhaps due to its film industry and milieu generating a wealth of content for viewers. The United States and United Kingdom figure prominently in watch hours, and Nigerian content is also popular in Cameroon and Ghana (in fact, it features in the top eight of all countries in the sample). The share of watch hours spent on US content is similar to that found in OECD countries; for instance, Cameroon is comparable to Mexico and Portugal, while the others are comparable to Denmark, Estonia, Norway, Sweden and the United Kingdom, whose shares of watch hours spent on US content are in the area of 34–39 percent.

	Share of Watch Hours Spent on Local Content, 2014 (%)	Increase in Share of Watch Hours Spent on Local Content, 2010–2014 (percentage points)	Top Three Content Countries, by Watch Hour Share	Share of Watch Hours Spent on US Content, 2014 (%)	Concentration of Watch Hours — Number of Countries Accounting for 80% of Watch Hours
Cameroon	3.14	1.46	United States, France, Nigeria	25.23	14
Ghana	9.92	4.56	United States, Nigeria, United Kingdom	35.87	10
Kenya	13.59	6.88	United States, Kenya, United Kingdom	36.69	12
Malawi	2.01	1.24	United States, United Kingdom, India	38.15	14
Nigeria	25.05	19.88	United States, Nigeria, United Kingdom	32.69	7
Rwanda	9.35	5.53	United States, Rwanda, France	34.10	14
Tanzania	13.64	9.76	United States, Tanzania, United Kingdom	32.37	12

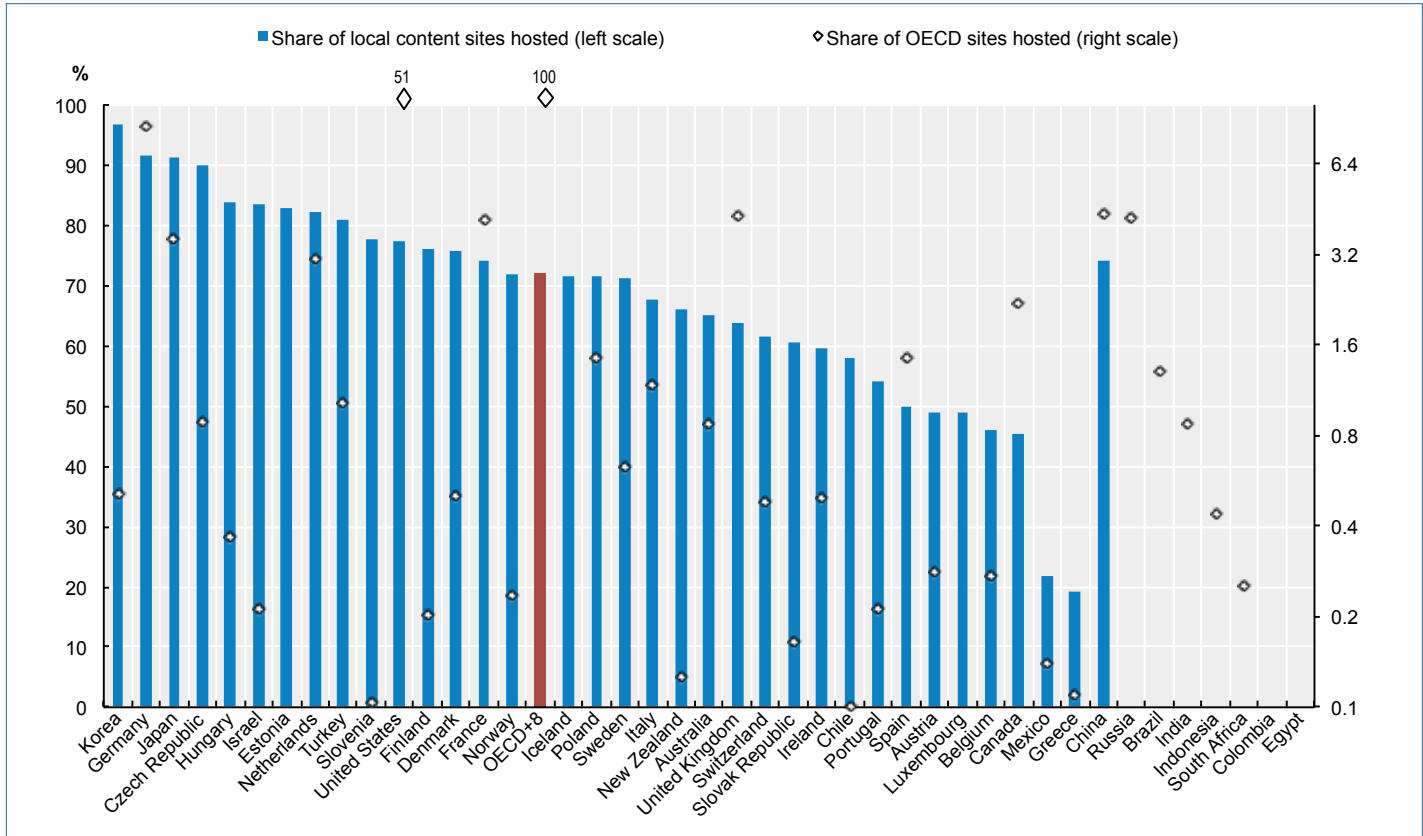
While the absolute number of watch hours is unknown, the 2010–2014 period was likely one of strong growth due to greater infrastructure provision. For instance, in 2009 there were no undersea cables connecting East Africa to the Internet, and only one cable serving the west and southern coasts. By 2013, numerous cables had been laid and some coastal countries are now served by multiple cables.¹ In-country infrastructure has also improved. There are now 37 Internet exchange points (IXPs) on the African continent (Packet Clearing House 2015) and at least two projects aim to advance regional and cross-border interconnection (AXIS and African Peering and Interconnection Forum).² Foreign companies are contributing — for example, in 2011 Google Global Cache³ was made available via the Kenyan IXP. Reductions in costs and latency significantly improved the user experience for video streaming (including YouTube) and Kenyans were able to more easily consume more local content, such as Kenyan news channels and TV programs. Local provider KENET reported a 10-fold increase in Google usage after the cache was created.

1 See maps developed by Steve Song at <https://manypossibilities.net/african-undersea-cables-a-history/>.

2 See <http://pages.au.int/axis/about> and www.internetsociety.org/events/afpif.

3 Google Global Cache is part of Google’s content delivery system, whereby Google servers are placed inside the network of network providers and ISPs to serve popular Google content, including YouTube.

Source: OECD calculations, based on information from source 4; Emily Taylor.

Figure 6: Local Content Sites Hosted in Country (2013)

Source: OECD (2015b), based on Pingdom, Alexa, and datacentermap.com

Note: Based on the analysis of 429,000 ccTLDs of the top one million sites. The remaining sites, including the gTLDs, were omitted from the list, as there is no reliable public data as to where the domains are registered. Data on the share of local content sites hosted was not available for Brazil, Colombia, Egypt, India, Indonesia, Russia and South Africa.

The location of hosting appears to go hand-in-hand with access to efficient infrastructure. Figure 6 shows that the United States accounts for a large share of the offshore market for hosting — it hosts 51 percent of all top sites in the OECD plus Brazil, China, Colombia, Egypt, India, Indonesia, Russia and South Africa. Figure 7 reveals a clear correlation between the number of co-location data centres¹³ and the number of top sites hosted in a country, suggesting that the favourable environment in the United States for setting up data centres (backhaul infrastructure, cost of energy/electricity, cost of land, regulatory environment) is an important factor in its pre-eminence. Germany is another popular location for hosting, along with France and the United Kingdom.


Logically, top hosting countries will be key conduits for data flows. For some businesses, there is a clear cost and efficiency advantage in routing data and content to data

centres in these locations. Aggregating data processing, for example, can enable better control over data practices, maximize the utilization of skilled staff and improve operational efficiency. Placing this activity in the most cost-efficient location is the best business choice.

Nevertheless, for some businesses there are advantages to keeping data and content close to consumers, not all of whom are in the top hosting locations. Growth in use of CDNs and caching of content close to customers are contributing to what is, in effect, economically driven localization of some data flows. Dennis Weller and Bill Woodcock (2013) note that CDN services, such as those provided by Akamai, have supported the demand for activities such as video streaming and downloading, while some large service providers, including Google, are building their own alternatives to transit (i.e., data centres). They note that where one end of a traffic flow is a server, especially a server holding non-unique information, then the data can be replicated in many locations in order to be closer to users.

This kind of structural change in the market makes routing more direct (thus reducing costs), improves quality and

13 The OECD (2014) identifies three types of data centres: in-house data centres, located with their organization; third-party data centres or co-location facilities that offer space to clients and compete on location (sites are often around large cities, capitals and financial centres), interconnection and energy efficiency; and Internet industry data centres, say, Amazon or Facebook, for which energy and land costs are crucial.

Figure 7: Co-location Data Centres and Top Sites Hosted

Source: OECD (2015b), based on Pingdom, Alexa and datacentermap.com.

Note: Number of top sites hosted based on the analysis of 429,000 ccTLDs of the top one million sites collected in 2013. The remaining sites, including the gTLDs, were omitted from the list, as there is no reliable public data as to where the domains are registered.

increases speed of delivery. But it also makes the analysis of cross-border data flows more complex, since what may once have been multiple cross-border flows of content (for example, a music video) can become one initial cross-border flow followed by multiple local downloads from a local cache. Internet openness remains important for the content to be shared, but the magnitude of content consumption enabled by that openness is less obviously seen in cross-border data flow data.

A key piece of shared infrastructure that enables data flows to stay local when economically logical is IXPs. IXPs enable the exchange of traffic via peering between connected networks, and their global distribution plays an important role in data flow routing. Crucially, the denser their presence, the more likely it is that data can flow across shorter and faster paths between its source and the destination. An analogy is with transport networks — must travellers transit through a distant hub or can they get to their destination more directly? The shorter the distance between customers and their IXP, the lower the costs and higher the quality of data flows.

Countries with a low density of IXPs are more likely to have cross-border data flows associated with their Internet activity, partly because IXPs and data centres are often co-

located,¹⁴ and partly because even if it involves a locally hosted site, data may have no choice but to transit through an IXP in another country to gain access to the destination network. Over time, the number of IXPs has grown, particularly in emerging economies. In April 2011, Weller and Woodcock (2013, 54) counted 357 IXPs worldwide, with 25 percent in North America and 38 percent in Europe. Prior to 2011, all regions had built new IXPs, with growth especially high in Latin America, which went from 20 to 34 IXPs. This growth was welcomed by the authors, as it reduced the need to “trombone” traffic out of the country or region, allowed for more direct routing of traffic and thus improved service quality, and freed up long-haul capacity to focus on actual out-of-region traffic (*ibid.*, 9).

As of October 2015, the global number of IXPs had grown to 452, with 60 in Latin America and 37 in Africa.¹⁵ The impetus to build an IXP essentially comes down to cost — ISPs prefer to have an IXP in close proximity so that

¹⁴ The OECD (2014) notes that carrier-neutral data centres endeavour to get IXPs into their facilities, as this makes interconnection with many networks possible.

¹⁵ Data obtained from https://prefix.pch.net/applications/ixpdir/menu_download.php.

the cost of outbound traffic is reduced.¹⁶ The break-even point depends on traffic volume and the ratio of local to international traffic — but at a cost of US\$3.50 per unit of megabits per second (Mbps) for IP transit, an ISP could be better off joining an IXP with a traffic volume of just 2,000 Mbps.¹⁷

WHERE TO FROM HERE?

The clear takeaway is that the flow of data across the Internet is complicated — data flows come in different forms, and they do not follow political or geographic borders but, rather, economic parameters that are set by changing market conditions and the regulatory/competitive environment. How, then, can we most usefully measure Internet openness so as to link it to indicators of governments' ultimate economic policy goals?

Looking ahead, two complementary approaches could be proposed as future research paths.

Approach one: Construct a global data flow data set that more accurately tracks geographical start and end points, as well as important waypoints en route, ideally with information on the types of flows, so as to better understand the nature and volume of data flows. This approach would essentially seek to build a data flow data set that could more easily be married with economic data sets, which are typically organized by country. Possible additional data and information sources to assist with this include:

- actual traffic data, both aggregate and in certain subcategories;
- further flow data from firms;
- information on the location of .com sites;
- information on the location of key data centre sites and their throughput; and
- information on barriers to data flows, to be used in constructing proxies for modelling purposes.

¹⁶ Weller and Woodcock (2013, Annex 4) describe how peering agreements, which comprise over 99 percent of all traffic exchange agreements, are constructed on the basis of equitable cost-revenue sharing between partners. This construction, in turn, relies on a distribution of IXPs that allows ISPs to have a similar balance of short- and long-haul paths to their traffic partners, so that neither is bearing disproportionately high costs.

¹⁷ The Internet Society (2014, 23) shows an example where traffic is destined for local termination and is either “local,” “near” or “far” from the IXP. Co-location costs are estimated at US\$1,000 per month, peering fees at US\$2,000 per month, equipment costs at US\$2,000 per month and transport into the IXP from US\$2,000–\$6,000 per month (depending on the distance). With an IP transit cost of US\$3.50 per Mbps (estimated from information from ISPs), the break-even point to join the IXP ranges from 2,000 to around 3,140 Mbps.

This approach raises the question of whether governments should seek to establish voluntary national statistical collections of traffic data. Australia, for instance, conducts a twice-yearly survey of ISPs with more than 1,000 subscribers, collecting data on *inter alia* the number of ISPs, subscriber sectors and the volume of data downloaded.¹⁸ It is perhaps time to explore whether such surveys should be expanded to include information on cross-border data flows. At the least, establishing a consistent cross-country methodology for collection of ISP data could enable analysis using domestic network traffic as a proxy for Internet openness, with coverage eventually expanding to cross-border data flows.

Approach two: Identify hotspots of data flow intensity (and, where possible, identify hotspots of data flow value) and overlay these with data showing the intensity and value of various economic performance variables (related to trade, innovation, entrepreneurship, productivity, and so on). In some ways, this approach would cast data flows as global data chains — similar to GVCs in the trade and production space — with intensity and value varying across different parts of the chain. Possible additional data and information sources to assist with this include:

- density of data infrastructure¹⁹: density and composition of players at IXPs; density of interconnection agreements at IXPs; bandwidth at IXPs; IP version 6 deployment by region; and
- analysis of value added of certain Internet-related activities, similar to analysis of trade in value added.

Finally, despite the evident need for further work, there are two important conclusions regarding Internet openness that emerge from this initial analysis. First, in line with its original design, the Internet remains a highly interdependent system. Data flows frequently have international dimensions and are not necessarily predictable. Reducing openness in any part of the system could have knock-on effects across the whole system, and thus all countries have an interest in ensuring that policy decisions regarding the Internet take into account the costs and benefits of openness. Given the important role of the United States in many aspects of the digital economy, its policy decisions clearly matter, but so too do those of other countries. For instance, Figure 6 showed Germany hosted almost 8.5 percent of OECD top websites in 2013, which suggests that its policy decisions on data flows and Internet openness would likely have significant consequences across the system.

¹⁸ See the Australian Bureau of Statistics Internet Activity Survey (catalogue 8153.0), www.abs.gov.au/ausstats/abs@.nsf/mf/8153.0/.

¹⁹ Presentation of this kind of data could take inspiration from the Internet connectivity maps produced by Larry Landweber in the 1990s. See <http://internethalloffame.org/news/in-their-own-words/larry-landweber-play-lab-world>.

Second, Internet openness, in terms of enabling data, information and knowledge to flow across the globe, is incontrovertibly tied to open markets and competitive conditions. Firms must be able to invest in or establish access to infrastructure that allows them to efficiently and effectively provide their services, be it on a local or cross-border basis; if they cannot, customer access, choice and service quality suffer. Weller and Woodcock (2013, 45) note a frequent observation that “improvement of the Internet depends upon a circular path of improvement of each component of the Internet’s infrastructure: IXPs, international connectivity, content, backbone networks, and access networks. One circumnavigates this circle endlessly, upgrading each in turn.” This observation has distinct parallels with Internet openness and suggests that measures of Internet openness need to incorporate infrastructural factors.

CONCLUSION

The initial stages of the OECD’s work to measure global data flows underscore the highly interconnected nature of today’s Internet. Aggregate information on Google searches and YouTube watch hours suggest that users are increasingly accessing content outside their countries, highlighting the potential of the Internet for cross-country information flows and knowledge exchange. In addition, countries’ Internet infrastructure and content have strong global interlinkages — one example being offshore hosting of local websites.

It is in the interests of all governments to improve the evidence base for policy making, because choices about Internet openness matter for countries’ trade and innovation performance. The strong international linkages inherent in the Internet also mean that the effects of a country’s Internet-related policies can spill across its borders. The OECD will continue to work with its members and partners to better understand global data flows and their effects on economies and societies.

WORKS CITED

- Accenture. 2013. *Entrepreneurial Innovation: How to Unleash a Key Source of Growth and Jobs in the G20 Countries: Young Entrepreneurs’ Alliance Summit 2013*. Accenture. www.g20yea.com/en/wp-content/uploads/Accenture-Entrepreneurial-Innovation-G20-YEA-Report.pdf.
- Blumenthal, Marjory S. and David D. Clark. 2001. “Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World.” *ACM Transactions on Internet Technology* 1 (1): 70–109.
- BSA. 2014. *Powering the Digital Economy: A Trade Agenda to Drive Growth*. Washington, DC: January. <http://digitaltrade.bsa.org/>.
- Cisco. 2015a. “The Zettabyte Era: Trends and Analysis.” White paper. May.
- . 2015b. “Cisco Visual Networking Index: Forecast and Methodology, 2014–2019.” White paper. May 27.
- Daigle, Leslie. 2015. *On the Nature of the Internet*. Global Commission on Internet Governance Paper Series No. 7. Waterloo, ON: CIGI and Chatham House. March 16. www.cigionline.org/publications/nature-of-internet.
- Dalberg. 2014. *Open for Business? The Economic Impact of Internet Openness*. Dalberg Global Development Advisors, March. www.dalberg.com/documents/Open_for_Business_Dalberg.pdf.
- Ezell, Stephen, Robert D. Atkinson and Michelle Wein. 2013. *Localization Barriers to Trade: Threat to the Global Innovation Economy*. Washington, DC: Information Technology and Innovation Foundation, September 25.
- Hofheinz, Paul and Michael Mandel. 2014. “Bridging the Data Gap: How Digital Innovation Can Drive Growth and Create Jobs.” Lisbon Council-Progressive Policy Institute Policy Brief 15/2014.
- . 2015. “Towards a 21st Century Agenda of Transatlantic Prosperity.” Lisbon Council-Progressive Policy Institute Policy Brief 19/2015.
- Internet Society. 2014. *The Internet Exchange Point Toolkit & Best Practices Guide: How to Maximize the Effectiveness of Independent Network Interconnection in Developing Regions and Emerging Markets*. Geneva, Switzerland and Virginia, United States: Internet Society February. www.internetsociety.org/sites/default/files/Global%20IXPToolkit_Collaborative%20Draft_Feb%202014.pdf.
- Kaplan, James M. and Kayvan Rowshankish. 2015. *Addressing the Impact of Data Location Regulation in Financial Services*. Global Commission on Internet Governance Paper Series No. 14. Waterloo, ON: CIGI and Chatham House. May 22. www.cigionline.org/publications/addressing-impact-of-data-location-regulation-financial-services.

- McKinsey Global Institute. 2014. *Global Flows in a Digital Age: How Trade, Finance, People, and Data Connect the World Economy*. McKinsey & Company, April.
- McNeil, Donald G., Jr. 2013. "Car Mechanic Dreams Up a Tool to Ease Births." *The New York Times*, November 13. www.nytimes.com/2013/11/14/health/new-tool-to-ease-difficult-births-a-plastic-bag.html?_r=0.
- Meltzer, Joshua. 2015. "The Internet, Cross-Border Data Flows and International Trade." *Asia and the Pacific Policy Studies* 2 (1): 90–102.
- Mettler, Ann and Anthony D. Williams. 2011. "The Rise of the Micro-Multinational: How Freelancers and Technology-Savvy Start-Ups Are Driving Growth, Jobs and Innovation." Lisbon Council Policy Brief V (3). Brussels, Belgium.
- . 2012. "Wired for Growth and Innovation: How Digital Technologies are Reshaping Small- and Medium-Sized Businesses." Lisbon Council Interactive Policy Brief 12/2012. Brussels, Belgium.
- Nicholson, Jessica R. and Ryan Noonan. 2014. "Digital Economy and Cross-Border Trade: The Value of Digitally-Deliverable Services." ESA Issue Brief #01-14. United States Department of Commerce: Economics and Statistics Administration, January 27.
- OECD. 2008. *OECD Guidelines for Multinational Enterprises*. Paris, France: OECD Publishing.
- . 2012. "The Impact of Internet in OECD Countries." OECD Digital Economy Papers, No. 200. Paris, France: OECD Publishing. <http://dx.doi.org/10.1787/5k962hhgpb5d-en>.
- . 2014. "International Cables, Gateways, Backhaul and International Exchange Points." OECD Digital Economy Papers, No. 232. Paris, France: OECD Publishing. <http://dx.doi.org/10.1787/5jz8m9jf3wkl-en>.
- . 2015a. *OECD Digital Economy Outlook 2015*. Paris, France: OECD Publishing. <http://dx.doi.org/10.1787/9789264232440-en>.
- . 2015b. *Data-Driven Innovation: Big Data for Growth and Well-Being*. Paris, France: OECD Publishing. <http://dx.doi.org/10.1787/9789264229358-en>.
- . Forthcoming 2016. *The Economic and Social Benefits of Internet Openness*. Paris, France: OECD Publishing.
- Packet Clearing House. 2015. "Packet Clearing House Report on Internet Exchange Point Locations." www.pch.net/ixpdir/summary.
- Swedish National Board of Trade. 2015. *No Transfer, No Production: A Report on Cross-border Data Transfers, Global Value Chains, and the Production of Goods*. Kommerskollegium.
- Weller, Dennis and Bill Woodcock. 2013. "Internet Traffic Exchange: Market Developments and Policy Challenges." OECD Digital Economy Papers, No. 207. OECD Publishing. <http://dx.doi.org/10.1787/5k918gpt130q-en>.
- West, Jeremy. 2016. *A Framework for Understanding Internet Openness*. Global Commission on Internet Governance Paper Series No. 35. Waterloo, ON: CIGI and Chatham House.
- Zwillenberg, Paul, Dominic Field and David Dean. 2014. *The Connected World: Greasing the Wheels of the Internet Economy: A Country-by-Country e-Friction Analysis*. Boston Consulting Group. www.bcgperspectives.com/content/articles/digital_economy_telecommunications_greasing_wheels_internet_economy/.

ABOUT THE AUTHOR

Sarah Box is counsellor to the directors in the Organisation for Economic Co-operation and Development (OECD) Directorate for Science, Technology and Innovation (DSTI). At the time of writing she was working in the Digital Economy Policy Division on issues of Internet openness. Since starting in DSTI in 2007 Sarah has worked on a variety of issues including human resources for science and technology, the OECD Innovation Strategy, public research organizations and the shipbuilding industry. Prior to joining the OECD, she worked as a senior research economist at the Australian Productivity Commission, and as an economist at the New Zealand Treasury

CHAPTER FOUR: TRACING THE ECONOMIC IMPACT OF REGULATIONS ON THE FREE FLOW OF DATA AND DATA LOCALIZATION

Matthias Bauer, Martina F. Ferracane and Erik van der Marel

Copyright © 2016 by Matthias Bauer, Martina F. Ferracane and Erik van der Marel

INTRODUCTION

Cross-border data regulation is a new type of regulation, which can impose significant costs on domestic and foreign firms (Christensen et al. 2013). There is, however, relatively little knowledge on the channels through which these data flow regulations affect the performance of the wider economy. To the knowledge of the authors, virtually no empirical assessment has been performed regarding the way in which data regulations affect the output of the wider economy. This chapter presents an empirical approach to uncover the relationship between regulations in domestic and cross-border data and the performance of the domestic economy. In addition, the chapter discusses the current policy debate related to data localization and its associated regulations for these emerging economies.

Regulation of data flows represents a relatively new feature in the broader spectrum of services regulation. It concerns rules on how personal data is utilized and processed by firms in the interaction between consumers and producers, or between producers. Consumers can be exposed to the release of their personal data on numerous occasions — for example, while using credit cards for economic transactions — or during instances ranging from using social media to accessing health care services. In many cases, the consumer and producer are located in different geographical locations, which motivates the transfer of data domestically or across borders. For instance, although a consumer’s credit card banking service might be located in the same country as the consumer, transaction data made are often stored on a server somewhere else, or are further processed in the banking affiliate’s head offices elsewhere in the world. Data flow regulations aim to regulate this flow of data between parties or across countries.

As in all services sectors, policy makers’ challenge is to find the right balance between developing necessary regulations that are linked to a particular social objective (or negative externality) and implementing these regulations at minimum cost, in terms of economic welfare, so they do not create an unnecessary cost burden for firms (Sáez et al. 2015). Yet, new rules on the regulation of cross-border consumer data for producers could also have detrimental economic effects (see Bauer et al. 2014). This is because data services regulations have a side effect of restricting transactions between domestic and foreign-using operators, which in turn limits the efficient sourcing of data processing

activities.¹ More importantly, today data are used by all sorts of so-called data-using industries (for example, downstream industries) as part of their input structure for production. In fact, services sectors are the main users of data. Regulatory restrictions of data can therefore inhibit the performance of sectors such as financial or business services or even new technological sectors using platforms.

In particular, this chapter shows the negative cost impact of data regulations on domestic industry performance in a select set of countries — namely emerging economies — by developing a regulatory data index that serves as a proxy indicator for regulations in data. As in Erik van der Marel, Hosuk Lee-Makiyama, Matthias Bauer and Bert Vershelde (forthcoming 2016), this chapter first makes a comprehensive assessment of the different types of regulatory barriers currently existing in various Organisation for Economic Co-operation and Development (OECD) and emerging economies to create a benchmark base of regulations. It then augments this benchmark to assess the extent to which the actually observed set of regulations in data in the eight emerging economy countries has a significant impact on downstream sectors that use data. To undertake this exercise, it is assumed that more intense users of data-processing services will be hurt to a greater extent than firms, for which data only account for a small share of total input use.²

Downstream industries that use data or data-related services to a greater degree are usually more dependent on the extent to which the transfer of data is regulated and/or freed from unnecessary cost-increasing regulatory measures than other sectors. This approach has been taken into account with the goal of assessing how regulations in data affect the overall economy. The chapter develops a benchmark of currently existing regulatory administrative barriers and later adds on the regulations in data to this benchmark. In a second step, the effect of these administrative regulations are estimated, including the ones in data, on the economic performance of the downstream users in terms of total factor productivity (TFP). In other words, the chapter augments a regulation index with the recently proposed data protection measures of eight countries included, and computes the

1 Domestic and foreign operators can trade these personal data as inputs at arm’s length (i.e., cross-border) or in-house through various ways, depending on the sourcing strategy. Following Sebastien Miroudot, Ranier Lanz and Alexandros Ragoussis (2009), if the data is traded within the firm with a foreign country it refers to *offshoring*; if traded outside the firm within the same country it refers to *domestic outsourcing*; if traded outside the firm and with a foreign country it refers to *global or international outsourcing*. Normally, trade data is recorded both when trade takes place in-house (intra-firm), as well as arm’s length (inter-firm) across borders, as it does not make a difference between the two channels.

2 The empirical strategy is borrowed from van der Marel et al. (forthcoming 2016) and follows Jens Matthias Arnold, Beata Javorcik and Aaditya Mattoo (2011) and Arnold, Javorcik, Molly Liscomp and Mattoo (2012), in which the authors have developed a similar index that identifies the extent to which downstream goods producers are affected by deregulations in services sectors.

precise TFP impact for these countries by sector as part of a counterfactual analysis.³ Admittedly, this approach is indirect, but robust, and borrowed from Erik van der Marel, Hosuk Lee-Makiyama, Matthias Bauer and Bert Vershelde (forthcoming 2016).

The first version of this study was conducted in 2014. In the meantime, extensive research was carried out, covering more than 60 countries around the world. Updated data will be publicly released in a comprehensive database by mid-2016. The data reveal that data regulations in general and data localization requirements in particular can take different forms, according to stated objectives, and affect downstream industries in many different economically distorting ways. A brief discussion of prominent measures and recent trends is provided after the empirical part of this chapter.

ASSESSMENT STRATEGY OF THE COST OF REGULATION

This section explores how data regulation is systematically related to the performance of the economy, divided into industries and services. An identification strategy has been applied based on three of information: first, the extent to which data regulation is present across various countries; second, the performance of downstream industries in each of the countries available in the dataset; and third, a measure that links this data regulation index to downstream industry performance.

COST PRICE INCREASES OF DATA REGULATION FOR DOMESTIC FIRMS

Currently, no composite index or indicator exists that measures the extent to which data and/or data services are regulated. Therefore, this chapter relies on an indirect approach of taking a pre-existing measure in order to ascertain a rough proxy indicator. This proxy index should only roughly reflect the regulations regarding data currently in place in the select countries, as it is chosen according to the types of real regulatory measures prevalent in the data usage of the select economies. The estimated outcomes of this indirect approach are then used to add on the real data regulations currently in place in the eight countries, and their inhibiting effect on economic performance are estimated.

A two-step procedure was undertaken to develop the proxy index. First, the actual regulations regarding the use, processing and cross-border transfer of data were verified across the select group of countries that are currently considering or implementing a regulatory package of data measures. By doing this, the *type* of regulations these are

³ The reason for looking at productivity is because higher costs of input usage will translate into lower rates of efficient usage of inputs in a firm or industry's production function.

was examined. Put differently, this method verifies what types of regulatory measures related to data are actually observable at this moment across the selection of countries, which have either proposed or already implemented data regulations. The actual laws for each country related to these data regulations are listed in Table 1. Based on this assessment of current regulatory barriers, a rough proxy index of the existing data regulations was selected and used as a benchmark to assess the cost-effect of data regulations in the wider economy.

In order to select this rough proxy variable, a sub-indicator of the integrated structure of the OECD Product Market Regulation (PMR) in services was used (as explained in Koske et al. 2015). The relevant indicators are "regulatory and administrative opacity" and "administrative burdens on start-ups," which relate to the approximate measures selected for this study. Within this PMR composition scheme, the sub-indicator called "data-producing sectors" was chosen. These indicators measure as close as possible the kind of prevailing regulatory barriers in the usage or process of data prevalent in sectors listed in Table 2. By doing so, this study relied on the information available on the types of regulations and market structure in these data sectors in the selected countries, as shown in Table 1.⁴ As such, this stage does not try to develop an exact index that measures the extent to which countries really regulate data, but instead makes a close match between currently existing regulations regarding data and the existing regulatory indexes currently available, sorted by the type (or based on the typology) of regulatory measures that currently prevails in data. Later, the real policy regulations of the select countries are added on to this benchmark index to estimate the real costs. As such, this rough proxy will be augmented.⁵

⁴ The countries selected for analysis of the prevailing (or considered) regulatory barriers in data services are the European Union, Brazil, China, Vietnam, Indonesia, Korea, Russia and India.

⁵ Another way of looking at this procedure is to think of a benchmark approach in which a point of reference is constructed from where those conducting the study, in a later stage, measure the costs associated with the actual implementation of data regulations observed in the countries taken up in this chapter.

Table 1: Types of Regulatory Barriers in Data Services

Type of Restriction	Regulatory Measure	Outcome
Restrictions related to the foreign supply of data services	Is there a data localization requirement?	Yes/Limited/No
Restrictions related to internal productivity losses/administrative costs	Is there a strict consent requirement for the collection, storage or dissemination of personal data?	Yes/No
	Does the law provide users with the right to review their stored information?	Yes/No
	Does the law provide users with the right to be forgotten/deleted?	Yes/No
	Is a notification of breaches toward the government/user obligatory?	Toward government/user/ government and user
	Are data protection impact assessments obligatory?	Yes/No
	Is a data protection officer required?	Yes/No/Qualified Yes
	Are there administrative sanctions for non-compliance? How high?	Varies according to height of sanctions
	Does the government require easy access to companies' data?	Yes/No
	Are companies required to retain data for a fixed period of time?	Yes/No

Data source: Authors; European Centre for International Political Economy <http://ecipe.org/>.

Table 2: Selected Sectors Related to Data Services

North American Industry Classification System 6-digit Sector	Description
511140, 511190	Directory, mailing list and other publishers
511210	Software publishers
516000	Internet publishing and broadcasting
517000	Telecommunications
518100	Internet Service Providers and Web search portals
518200	Data processing, hosting and related services
519000	Other information services
541511	Custom computer programming services
541512	Computer system design services
541513, 541519	Other computer-related services, including facilities management

Data source: BEA www.bea.gov/industry/io_annual.htm.

For the index, administrative regulations were selected, which in the PMR falls under the division “barriers to entrepreneurship” and is made up of a simple unweighted average between two indicators measuring administrative barriers: “regulatory and administrative opacity” and “administrative burdens on start-ups.” Of course, these two categories do not exactly match data regulations, but since most of the data regulations are administrative in nature, this index was used to be as close as possible for the benchmark needed for this study. In other words, most of the regulatory barriers related to data observed in

sectors (listed in Table 1) are of an administrative character, which is what these PMR sub-indexes appear to measure. So by selecting the administrative barriers, this study tries to capture at least some of the regulatory burdens that are also likely to prevail in data.⁶

As a regulatory index of data or data services was not developed over this more indirect approach, this chapter tries to correct for bias as much as possible by multiplying the countrywide proxy index with the share of data services used for each sector — i.e., the so-called industry usage of data.

DOWNSTREAM LINKAGE

How does one address the link between this proxy index of regulation to each downstream sector using data in their input production so that the potential cost of data regulation can be measured for the wider economy? Note that an unweighted approach, in which the regulatory barriers are directly linked to each downstream industry, would be insufficient to properly capture the real economic effects of the performance variable. This is because some sectors are more dependent on data than others. Hence, to finalize the benchmark approach, the proxy index of administrative regulation was connected to each individual data-using sector in the economy before measuring its impact on the economy-wide output performance of each downstream industry.

⁶ In van der Marel et al. (forthcoming 2016), other sub-indicators were checked to see if they provided different outcomes, as a robustness check.

This calculates the data intensity for each downstream using sector of data in a typical economy using US Input-Output Use tables from the U.S. Bureau of Economic Analysis (BEA). Through this, the share of data-reliance for each industry and sector for a representative economy was computed. One advantage of taking these Input-Output Use matrixes is their level of disaggregation. More than 400 sectors are specified at six-digit commodity and services code level, which makes it the most detailed input-output table currently available across the globe. This weighted approach was selected because regulation in data will be most felt in industries and services sectors that use data and data-related services most intensively as inputs for the production process of other goods and services. Put differently, the input range of the data and data-related activities is likely to be more important for some manufacturing and services industries than others, and this variability is exploited in the cross-sectional panel's empirical setting.⁷

Table 3 shows this share of data reliance, which is the proportion of input used by each sector from the data-producing sectors listed in Table 2. This list is largely based on the information technology (IT)-producing sectors put forward in Dale W. Jorgenson, Mun S. Ho and Jon D. Samuels (2005; 2007; 2010). The only difference is that sectors not included in this chapter's selection are the IT equipment manufacturing sectors, which are pure goods industries that deal with the transfer of data to a much lower degree.⁸ With the distinction between data producers and users in mind, the data intensity of inputs provided by the data producers for each downstream industry can be calculated. Table 3 shows that telecommunications, information and communication technology (ICT) business services, finance and insurance are sectors that use data services most intensely, whereas the processed foods industry, metals industry and primary agriculture are sectors where data services play only a negligible role. Overall, data are used much more in services than in goods industries.

7 This approach follows the strategy taken in Arnold et al. (2011; 2012), in which the authors developed an index that identifies the extent to which downstream producers are affected by regulations in services sectors. In their seminal studies, downstream industries using services to a greater extent are considered more dependent on the degree to which services are liberalized or freed from cost-increasing regulatory measures. This section applies their dependency index, but then only for the usage of data.

8 These industries are the following: computer and peripheral equipment manufacturing; communications equipment manufacturing; and semiconductor and other electronic component manufacturing (Jorgensen, Ho and Samuels 2010). As the authors of this chapter see it, cross-border data is a new emerging phenomenon that closely resembles services, as the types of regulatory barriers found in data are extremely close to those found in services sectors, hence, the interchangeable usage of data and data services and data-processing services. The selection of sectors in Table 1 was done by an expert group that worked closely with data services companies and representatives.

Table 3: Data Intensities

GTAP Sector	Sector Description	Data Intensity
communication	Post and telecommunication services	0.318
obsict	Other business and ICT services	0.069
fininsurance	Financial and insurance services	0.050
machinery	Machinery and electronic equipment	0.049
oconsumer	Other consumer services	0.048
oservices	Public services, dwellings	0.040
distribution	Trade and distribution services	0.037
water	Water and other utility services	0.034
transport	Transport services	0.032
construction	Construction	0.024
othermanuf	Manufactures nec.	0.024
fabmetals	Metal products	0.020
nonmetmin	Mineral products nec.	0.014
lumberpaper	Wood and paper products	0.014
energy	Coal, petroleum and gas production	0.011
transequip	Motor vehicles and parts	0.008
chemicals	Chemicals, rubber and plastic products	0.008
bevtextcloth	Beverages/tobacco products; clothing and leather products	0.007
metals	Ferrous metals and metals nec.	0.007
primagrother	Primary agricultural products	0.007
procfoods	Meat, vegetable oils, dairy, sugar and food products nec.	0.006

Data source: Author's calculations using BEA at www.bea.gov/industry/io_annual.htm.

Note: nec. = not elsewhere classified.

Having computed the types of information of administrative regulatory measures for each country and data services intensities for each downstream sector, these two data variables are linked into one indicator to obtain the following weighted index for the so-called data regulation linkage (DRL),

$$DRL_{oit} = \sum_k \alpha_{ik} \text{regulation index}_{ot} \quad (1)$$

where DRL stands for the data regulation linkage for sector i in a typical country o in year t , which is measured by the proportion α_{ik} of inputs sourced by both the manufacturing and services sectors i from data services sectors k , multiplied with the proxy index for data services regulation for each country o in year t . Again, the variable α_{ik} is sector-specific and calculated using the BEA's US Input-Output Use tables as presented in Table 3, while the *regulation index*_{ot} variable is computed using the administrative barriers index from the PMR structure as previously explained. Hence, according to

equation (1) the input coefficients in terms of data intensities presented in Table 3 are multiplied with the PMR regulation index numbers.⁹

MEASURING THE PERFORMANCE OF THE ECONOMY

The final stage of this first step is to measure the extent to which the benchmark of administrative regulations for data (i.e., the DRL) has an effect on the performance of the whole economy. To take stock of the economy-wide output performance, two different variables were studied, which are inversely related to each other, namely the TFP and a price index based on value added calculations (Pva). The reason for selecting TFP as an output performance measure follows standard practise in the international economic literature. More regulations will inhibit firms from sourcing input efficiently, which will show up in higher costs for firms and industries as part of their production function. This, in turn, will increase prices, which will eventually translate to lower productivity, i.e., TFP.

As an example, when data localization is put in force it means companies are prohibited from sending data back and forth freely between affiliates or companies by adhering to strict rules of local storage or other administrative regulations. This increases costs for firms to source and process data efficiently, which will increase business operations' complexity and eventually decrease productivity. Another example is data protection impact assessment implemented by countries. This regulatory barrier will affect firms working with lots of data more than other less data-related firms, which could lower their relative efficient production. In order to find out a meaningful conclusion of the impact of data regulations on downstream TFP and price performance, standard parametric estimations techniques are used. The following estimation equations are used:

$$\ln(TFP)_{oit} = \alpha_i + \beta_1 DRL_{oit} + \gamma_o + \delta_i + \zeta_t + \varepsilon_{oit} \quad (2a)$$

$$\ln(Pva)_{oit} = \alpha_i + \beta_1 DRL_{oit} + \gamma_o + \delta_i + \zeta_t + \varepsilon_{oit} \quad (2b)$$

where TFP and Pva in industry i in country o in year t is explained by the data regulation linkage index for that same industry i in country o in year t in both equation (2a) and (2b) and are put in logs. In both equations, the terms γ_o , δ_i , and ζ_t stand for the fixed effect by country, sector and year respectively, which are also included in the empirical model. These fixed effects take care of the issue that other factors undoubtedly may also have an effect on TFP. For

9 Note that the group of countries over which the DRL is calculated spans a wider selection than the list of countries mentioned in footnote 4. This is because these countries are used as examples of governments where data regulations have recently been put forward that have served for selecting a close proxy of similar types of regulations. The availability of these proxy regulations from the PMR, as explained in the section on "Quantification Process," covers almost all OECD and emerging economies.

instance, the extent to which a country accumulates high-skilled labour or ICT-related capital could also affect productivity. Since the fixed effect picks up this variation by country and sector, no additional control variables will have to be included. Data for both TFP and prices are taken from the EU KLEMS database, which covers information for two-digit sectors based on the NACE classification and are calculated on a value-added basis. Finally, both equations' error term is given by ε_{oit} . Altogether, there is a small panel dataset for three years covering 21 goods and services sectors for 12 countries. Summary statistics of all these variables are given in Table 4.

Table 4: Summary Statistics for Dataset

Index Proxy	Period	Mean	Std. Dev.	Min.	Max.	No. of Observations
TFP (logs)	All	4.566	0.171	3.292	6.338	996
Pva (logs)	All	4.590	0.201	2.054	5.768	1002
DRL	All	0.084	0.136	0.004	0.965	1008

Source: Authors.

Table 5 provides the results of the regressions. The coefficients have the expected negative sign on TFP and positive sign on value-added prices. Both coefficient results for TFP and prices are statistically significant. The results suggest that administrative regulatory barriers in sectors using data-processing services most intensively exhibit a dampening effect on TFP, while also exerting an upward pressure on prices in these sectors. A one standard deviation change in the DRL variable would therefore decrease TFP on average by 3.9 percent. Similarly, for prices, a one standard-deviation change in the DRL would increase prices on average by 5.3 percent.¹⁰

Table 5: Regression Results on Prices and TFP

	(1)	(2)
	lnTFP	lnPRICE
DRL	-0.255**	0.395***
	(0.122)	(0.108)
Observations	996	1,002
R-squared	0.159	0.173
RMSE	0.164	0.187

Note: Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1
Source: Authors.

10 Van der Marel et al. (forthcoming 2016) provide further regressions output results and take stock of any endogeneity concern (which is not included in this chapter). Rather than assuming regulatory policy in data services affects downstream TFP in the wider economy, taking stock of this potential endogeneity means the exclusion of any reversed causality — i.e., firms that already perform well in terms of TFP are the ones lobbying for precisely lower regulatory barriers. Taking the lag on the independent variables shifts the time frame backwards so this possibility of lobbying is ruled out, as firms cannot influence policy that took place in previous years.

APPLICATION OF COST PRICE INCREASES OF DATA REGULATION

Based on the econometric exercise, the second step performs a counterfactual analysis for various countries that have, in reality, implemented a package of regulatory measures related to data. For this chapter's research purposes, various emerging market economies, plus South Korea and the European Union, were selected for the counterfactual analysis to have a small variety of different income countries where data regulations are currently observable. This is done by taking into consideration the data-processing services regulation laws currently under consideration as listed in Table 6. As a result, the elasticity results from other benchmark approaches in Table 5 are used to estimate the TFP losses associated with the actual implementation of data regulations observable in these eight countries.

Table 6: Selected Economies with Law Proposals for Data Processing

Country	Title Law for Data-processing Barriers
Brazil	Marco Civil
China	Decision on Strengthening the Protection of Information on the Internet (December 2012) and Telecommunication and Internet User's Personal Information Protection Measures (September 2013), plus Consumer Protection Law amendment of April 2013
India	Data retention provision of Information Technology Act, proposed National Security Council Secretariat strategy on cyber security plus proposed licensing requirement by Department of Telecom
Indonesia	GR 98 (2012) and EIT Law (2008)
South Korea	Personal Information Protection Act
European Union	EU General Data Protection Regulation
Vietnam	Decree 72
Russia	Federal Law No. 152-FZ and Federal Law no. 242-FZ

Source: Authors.

QUANTIFICATION PROCESS

Whereas the original index of administrative barriers was predefined and constructed as part of the OECD's PMR database, this time the index was augmented with actual observed administrative barriers in data and/or data services in the selected countries. Thus, in the quantification process, the de facto implemented regulatory barriers are added on top of the existing regulatory barriers used in the previous analysis. This is done by analyzing and quantifying the proposed data regulation laws currently in deliberation in the aforementioned countries (as presented in Table 6) in the same way as the original index of administrative barriers as part of the PMR structure.

This is done by selecting those data-related regulatory barriers that first, the selected countries have in common across their proposed law programs; and second, are likely to have a significant cost burden for firms. This selection has been carried out in close contact with experts in the field. The final selection of barriers is presented in Table 1. This selection process aims to include only those measures that have an economically important weight in terms of additional costs for the firm, as previously explained. Therefore, in order to assess whether these various barriers would really carry along significant costs for data services providers, various secondary sources were relied upon (see Christensen et al. 2013; Le Merle et al. 2012; UK Ministry of Justice 2012; PriceWaterhouseCoopers 2013; European Commission 2012). These sources give information about the excessive cost burden of the various regulatory data policies. With this selection of the regulatory data barriers at hand, appropriate weights are applied for each regulatory measure to take into account the average relative importance of each of the policy measures based on the expert judgements approach and based on these reports. The results are shown in Table 7.¹¹

Eventually, based on this coding scheme and the application of weights, a new index was derived that ranges between zero and six for each of the selected countries, which is consistent with the PMR score. The results for each country are given in Table 8. The final score is computed in the last row of the table. This measure indicates that a higher score for this index means that countries have implemented a greater degree of regulations in data services. The highest index can be found for Russia (4.82), followed by China (3.88) and Korea (3.82). Brazil (0.75), Vietnam (2.19) and India (2.36) have relatively low index levels of data restrictions. Note that this is due in large part to putting a higher weight on the barrier of data localization in this methodology. Having this type of barrier in place explains the relatively high score for Russia and China, whereas the European Union has a relatively high score because of many other domestic administrative barriers in place.

¹¹ See van der Marel et al. (forthcoming 2016) for further discussion on this issue. The discussion of data localization in various countries shows that definitions of data privacy, personal data and the obligation to store and disentangle certain categories of data creates various types of costs that can also go beyond pure administrative cost. These costs are comprised of, among others, business operation risks and the risk of additional security breaches due to external attacks by hackers. These costs also vary dramatically between those sectors that use data intensively, such as logistics and communication services, and those sectors that use data less intensively, such as primary sectors (excluding marketing and sales of commodities and raw materials).

Table 7: Quantification of Proposed Data-processing Barriers

	Weights by Theme (bj)	Question Weights (ck)	Coding of Outcome Data		
			No	Limited	Yes
Foreign supply of data services:	0.3		No	Limited	Yes
Is there a data localization requirement?		1	0	3	6
Internal administrative costs measures:	0.7				
			No		Yes
Is there a strict consent requirement for the collection, storage and dissemination of personal data?		0.050	0		6
Does the law provide users with the right to review their stored information?		0.050	0		6
Does the law provide users with the right to be forgotten, deleted?		0.047	0		6
			No	Government or user	Both
Is notification of breaches towards the government and/or users obligatory?		0.200	0	3	6
			No		Yes
Are data protection impact assessments obligatory?		0.175	0		6
Is a data protection officer required?		0.375	0		6
			No	Some	High
Are there administrative sanctions for non-compliance? How much?		0.047	0	3	6
			No		Yes
Does the government require easy access to companies' data?		0.047	0		6
Are firms required to retain data for a fixed period of time?		0.013	0		6
Country scores (0–6)			$\sum_j (b_j) \sum_k (c_k) \text{ answer}_{jk}$		

Source: Authors.

Note: Question weights are based on Christensen et al. (2013) and UK Ministry of Justice (2012).

AUGMENTING THE INDEX FOR ADMINISTRATIVE BARRIERS

The next step is to augment the existing index of administrative barriers in services with the index created for administrative barriers in data services. This is done by distinguishing between two periods of time, namely one where these data services barriers have not been put in place ($t=0$) as per today and which have been taken up in the empirical exercise in previous sections, and a hypothetical time period in which the data-related laws are implemented and are now applied in $t+1$. In other words, the initial index is augmented with the index created in Table 8, which describes what countries have implemented in terms of new data regulatory measures.

In $t=0$, a weighted average of both administrative barriers indices are applied, as defined with the DRL in the section “Cost Price Increases of Data Regulation for Domestic Firms,” plus an assumed index set to zero for administrative

barriers related to data not implemented yet in time $t=0$, the current year. Both a 0.4 weight to the two administrative barriers and only a 0.2 to the third index of data regulations are applied, because data regulations affect the use of data, which is still a part of the economy relatively lower in size than the part of the economy of all other (services) sector for which the existing barriers are targeted.¹² In the following step, data regulations in Table 8 are now implemented in time period $t+1$ as the third indicator.

¹² Obviously, this weighting scheme is somewhat arbitrary in the sense that one could also think of a lower weight. Nonetheless, from a methodological point of view, this matters less as one could adjust this weighting scheme accordingly when estimating the cost impact of data regulations.

Table 8: Index Outcomes of the Quantification Method

	Russia	China	Korea	EU	Indonesia	India	Vietnam	Brazil
Foreign supply of data services:								
Is there a data localization requirement?	6	6	3	0	6	6	6	0
Internal administrative costs measures:								
Is there a strict consent requirement for the collection, storage and dissemination of personal data?	6	6	6	3	3	6	0	6
Does the law provide users with the right to review their stored information?	6	0	6	6	6	0	0	0
Does the law provide users with the right to be forgotten, deleted?	6	6	6	6	0	0	6	6
Is notification of breaches towards the government and/or users obligatory?	6	6	6	3	3	0	0	0
Are data protection impact assessments obligatory?	0	6	0	6	0	0	0	0
Is a data protection officer required?	6	0	6	6	0	0	0	0
Are there administrative sanctions for non-compliance? How much?	6	3	3	3	3	3	0	6
Does the government require easy access to companies' data?	0	0	0	0	0	6	6	3
Are firms required to retain data for a fixed period of time?	0	0	0	6	0	6	0	6
Country scores (0–6)	4.82	3.88	3.82	3.18	2.42	2.36	2.19	0.75

Source: Authors.

In other words, this time period computes the same weighted average of the administrative barriers as defined under the “Cost Price Increases” section, plus the index obtained in Table 8 instead of the assumed index of zero for the previous time period in which no new data regulations were yet implemented. Again, the same weighted average of these three indicators is applied just in the previous time period so as to take account of the size of the data services sector in the whole economy, which will prevent any overestimation of the indexes. The difference between these two time periods is that this index now measures in a scalable way the increased regulatory costs for countries that have implemented additional real regulations on data as part of their overall set of administrative barriers. Overall, a higher index means that a country has a more restrictive data services regulatory regime.

Table 9 shows the results for both time periods. Unsurprisingly, the highest increase in the index can be observed for China, Korea and the European Union. The lowest movement between the two time periods is found for Brazil, in large part because Brazil has not implemented

laws related to data localization in addition to some of the other barriers receiving high weights in the methodology.

Table 9: Index Movements between Period (t=0) and (t+1) for Augmented Index

	Index (t=0)	Index (t+1)
Brazil	0.58	0.63
China	0.78	1.04
India	0.86	1.02
Indonesia	0.24	0.40
Korea	0.21	0.47
EU	0.34	0.56
Vietnam	0.78	0.92
Russia	0.44	0.76

Source: Authors.

Note: Each time period contains an average of the administrative barriers as outlined in the “Quantification Process” section and the regulations related to data. In t=0 the index for data processing services is set to zero, whereas in t+1 the index for data processing services is set to the level as defined in Table 8 for each country.

CALCULATING TFP LOSSES

Finally, this study calculates potential TFP losses as a consequence of the counterfactual situation in which countries are in fact implementing their regulatory law programs on data. The indices in both time periods are used and these figures are plugged into the equation (2a) so that two different TFP levels are obtained: one before the implementation of data laws ($t=0$) and one after ($t+1$). TFP calculations using equation (2a) applies the coefficient results as found in Table 5 (i.e., the β_1), information on the data intensities as presented in Table 3 and the fixed effects by sector, which are acquired from running the regressions. After plugging in the data and calculating the $\ln(\text{TFP})$ for both periods, the percentage change in TFP from taking the first difference of $\ln(\text{TFP})_{t+1}$ and $\ln(\text{TFP})_{t=0}$ is obtained.

The results of the estimated TFP changes are presented in Table 10, which gives a situation of the downstream

productivity effects as a consequence of the implementation of the current data regulations under consideration or which have already been applied. The sectors are sorted by the size of the TFP reduction. As one can see, the communication sector experiences the greatest losses since the effect in this sector is most likely caused by its high dependency on data input use. For instance, in Korea, China and the European Union, the percentage of TFP losses are estimated to be around two percent in communications. Similarly, both the ICT business services sector and the finance and insurance sector also experience relatively high TFP losses of around 0.34 percent in China. Machinery is a merchandise sector in the ranking where a relatively high TFP losses would take place. At the bottom of the list are the other primary agricultural sector and the processed foods industry with only minor TFP losses. Their input dependency on data is likely to be very small. Overall, the results show that it is the services economy that will suffer most from regulations in data services.

Table 10: TFP Changes as a Consequence of Data-processing Regulations

Sector	Brazil	China	India	India	Korea	EU	Vietnam	Russia
All sectors	-0.07	-0.35	-0.22	-0.22	-0.35	-0.29	-0.20	-0.44
Goods	-0.02	-0.12	-0.07	-0.08	-0.12	-0.10	-0.07	-0.15
Services	-0.10	-0.52	-0.32	-0.32	-0.51	-0.43	-0.29	-0.64
Business services	-0.17	-0.85	-0.52	-0.53	-0.84	-0.70	-0.48	-1.06
communication	-0.42	-2.16	-1.31	-1.35	-2.13	-1.77	-1.22	-2.68
obsict	-0.09	-0.47	-0.29	-0.29	-0.46	-0.39	-0.27	-0.57
fininsurance	-0.07	-0.34	-0.21	-0.21	-0.34	-0.28	-0.19	-0.43
machinery	-0.07	-0.34	-0.20	-0.21	-0.33	-0.28	-0.19	-0.42
oconsumer	-0.06	-0.33	-0.20	-0.20	-0.32	-0.27	-0.18	-0.41
oservices	-0.05	-0.27	-0.17	-0.17	-0.27	-0.22	-0.15	-0.34
distribution	-0.05	-0.25	-0.15	-0.16	-0.25	-0.21	-0.14	-0.32
water	-0.04	-0.23	-0.14	-0.14	-0.23	-0.19	-0.13	-0.29
transport	-0.04	-0.22	-0.13	-0.14	-0.22	-0.18	-0.12	-0.27
construction	-0.03	-0.16	-0.10	-0.10	-0.16	-0.13	-0.09	-0.20
othermanuf	-0.03	-0.16	-0.10	-0.10	-0.16	-0.13	-0.09	-0.20
fabmetals	-0.03	-0.14	-0.08	-0.08	-0.13	-0.11	-0.08	-0.17
nonmetmin	-0.02	-0.10	-0.06	-0.06	-0.10	-0.08	-0.06	-0.12
lumberpaper	-0.02	-0.09	-0.06	-0.06	-0.09	-0.08	-0.05	-0.12
energy	-0.01	-0.07	-0.05	-0.05	-0.07	-0.06	-0.04	-0.09
transequip	-0.01	-0.06	-0.03	-0.04	-0.06	-0.05	-0.03	-0.07
chemicals	-0.01	-0.06	-0.03	-0.04	-0.06	-0.05	-0.03	-0.07
bevtextcloth	-0.01	-0.05	-0.03	-0.03	-0.05	-0.04	-0.03	-0.06
metals	-0.01	-0.05	-0.03	-0.03	-0.05	-0.04	-0.03	-0.06
primagrother	-0.01	-0.04	-0.03	-0.03	-0.04	-0.04	-0.03	-0.06
procfoods	-0.01	-0.04	-0.03	-0.03	-0.04	-0.03	-0.02	-0.05

Source: Authors.

Note: Sectors follow the GTAP classification.

SIMULATION OF DATA REGULATIONS' IMPACT ON ECONOMIC OUTPUT AND TRADE

The downstream TFP estimates are applied to a computable general equilibrium (CGE) model, in which the wider macroeconomic impact of regulations in data and the impact on industrial output and trade volumes are estimated. The model applied in this study is Global Trade and Analysis Project 8, a commonly applied CGE model in the international trade literature.¹³ The simulation results are presented in Tables 11–14. Everything else being equal, the simulations indicate that in the medium- to long-term losses in economic activity (real GDP) range from 0.10 percent for Brazil to 0.48 percent for the European Union, 0.55 percent for China and 0.58 percent for Korea. Both Korea and the European Union already have fairly strict data regulations at the baseline, i.e., the scenario for which the economic impact has been estimated. In addition, their economies are specializing in sectors that are relatively data-intensive, which largely explains their great losses.

¹³ A more detailed description of the model and the applied methodology is given in van der Marel et al. (forthcoming 2016). Note that the results of estimation merely have an indicative character. CGE models are not fit for forecasting the very precise macroeconomic impact of the regulatory barriers in data services to trade. However, CGE models give us an estimate about the direction of the results of this econometric exercise, which is why it is employed in this chapter.

Table 11: Simulation Results and Percentage Changes in Real GDP

Change in Real GDP	
EU28	-0.48
Brazil	-0.10
China	-0.55
India	-0.25
Indonesia	-0.23
Korea	-0.58
Vietnam	-0.24

Source: Authors' calculations.

As regards industrial output, the production of data-intensive manufacturing and services sectors shrinks in all countries, while less data-intensive sectors — such as agriculture, food and textiles — generally grow in absolute and relative terms. Losses are notably taking place in the services sectors, with greatest decrease of sectoral output in sectors such as communications and business services, as well as finance and construction. The estimated changes in the countries' production patterns are also reflected by percentage changes in industrial trade balances. The strongest sectoral impact is found for trade in communication and business services. Since domestic production of communication and business services becomes less competitive vis-à-vis foreign suppliers, exports drop.

Table 12: Simulation Results and Percentage Changes in Sectoral Production

		Grains & Crops	Meat & Lifestock	Extraction	Processed Food	Textiles	Manufacturing	Distribution	Utilities	Communication Services	Business Services	Financial Services	Water	Transport Services	Construction	Other Services
EU28	Scenario 1	0.16	-0.10	0.16	-0.15	0.23	-0.04	-0.10	-0.06	-0.74	-0.25	-0.22	-0.21	-0.28	-0.56	-0.37
	Scenario 2	0.24	-0.12	0.26	-0.19	0.37	-0.06	-0.15	-0.07	-1.13	-0.25	-0.31	-0.30	-0.39	-0.79	-0.55
	Scenario 3	0.29	-0.13	0.31	-0.23	0.45	-0.07	-0.18	-0.09	-1.36	-0.30	-0.37	-0.36	-0.47	-0.94	-0.65
Brazil	Scenario 1	0.06	0.03	0.04	-0.01	-0.01	-0.02	-0.02	-0.01	-0.04	-0.04	-0.07	-0.04	-0.05	-0.14	-0.10
	Scenario 2	0.22	0.09	0.14	-0.01	-0.03	-0.09	-0.06	-0.04	-0.15	-0.06	-0.23	-0.15	-0.18	-0.47	-0.35
	Scenario 3	0.50	0.21	0.32	-0.03	-0.08	-0.21	-0.14	-0.09	-0.35	-0.14	-0.53	-0.35	-0.4	-1.07	-0.82
China	Scenario 1	0.03	-0.11	0.04	-0.06	0.23	-0.24	-0.38	-0.08	-0.59	-0.33	-0.21	-0.12	-0.17	-0.38	-0.43
	Scenario 2	0.04	-0.11	0.04	-0.08	0.23	-0.24	-0.31	-0.08	-0.59	-0.33	-0.21	-0.13	-0.17	-0.38	-0.43
	Scenario 3	0.06	-0.17	0.07	-0.12	0.38	-0.38	-0.59	-0.13	-0.91	-0.51	-0.32	-0.20	-0.26	-0.60	-0.67
India	Scenario 1	0.00	-0.10	0.05	-0.07	0.10	-0.17	0.02	-0.05	-0.60	-0.30	-0.12	-0.16	-0.13	-0.17	-0.22
	Scenario 2	0.00	-0.10	0.05	-0.07	0.10	-0.17	0.02	-0.05	-0.60	-0.30	-0.12	-0.16	-0.13	-0.17	-0.22
	Scenario 3	0.01	-0.23	0.14	-0.15	0.30	-0.42	0.07	-0.11	-1.51	-0.70	-0.27	-0.37	-0.30	-0.44	-0.23
Indonesia	Scenario 1	0.01	-0.06	0.05	-0.01	0.24	-0.18	-0.04	-0.12	-0.97	-0.24	-0.16	-0.20	-0.13	-0.25	-0.21
	Scenario 2	0.01	-0.06	0.05	-0.01	0.25	-0.17	-0.03	-0.12	-0.97	-0.33	-0.16	-0.20	-0.13	-0.26	-0.22
	Scenario 3	0.04	-0.14	0.13	0.00	0.64	-0.44	-0.09	-0.29	-2.40	-0.83	-0.40	-0.50	-0.33	-0.64	-0.53
Korea	Scenario 1	0.12	-0.17	0.18	-0.24	0.43	-0.07	-0.65	-0.02	-0.83	-0.22	-0.24	-0.19	-0.29	-0.59	-0.45
	Scenario 2	0.14	-0.21	0.23	-0.30	0.55	-0.08	-0.81	-0.02	-1.03	-0.28	-0.30	-0.23	-0.36	-0.73	-0.55
	Scenario 3	0.16	-0.27	0.30	-0.38	0.71	-0.11	-1.02	-0.03	-1.31	-0.36	-0.38	-0.29	-0.46	-0.92	-0.70
Vietnam	Scenario 1	0.06	-0.11	0.03	0.02	0.08	-0.15	0.04	-0.09	-1.19	-0.16	-0.14	-0.01	-0.18	-0.29	-0.17
	Scenario 2	0.06	-0.11	0.03	0.02	0.08	-0.15	0.04	-0.09	-1.19	-0.16	-0.14	-0.01	-0.18	-0.29	-0.17
	Scenario 3	0.18	-0.30	0.11	0.06	0.31	-0.42	0.12	-0.26	-3.28	-0.42	-0.42	-0.02	-0.50	-0.79	-0.48

Source: Authors' calculations.

Table 13: Simulation Results and Percentage Changes in Sectoral Imports

		Grains & Crops	Meat and Lifestock	Extraction	Processed Food	Textiles	Manufacturing	Distribution	Utilities	Communication Services	Business Services	Financial Services	Water	Transport Services	Construction	Other Services
EU28	Scenario 1	-0.29	-0.38	-0.01	-0.35	-0.49	-0.25	-0.39	-0.26	0.93	-0.21	-0.06	-0.53	-0.26	-0.61	-0.51
	Scenario 2	-0.38	-0.52	0.00	-0.49	-0.69	-0.33	-0.51	-0.35	1.51	-0.40	-0.06	-0.73	-0.35	-0.85	-0.71
	Scenario 3	-0.47	-0.63	0.00	-0.59	-0.83	-0.39	-0.62	-0.42	1.82	-0.47	-0.06	-0.87	-0.42	-1.01	-0.85
Brazil	Scenario 1	-0.14	-0.24	-0.03	-0.17	-0.25	-0.10	0.06	-0.11	0.84	0.07	0.00	-0.07	-0.20	-0.21	-0.19
	Scenario 2	-0.49	-0.83	-0.10	-0.57	-0.85	-0.33	0.20	-0.36	2.82	0.04	-0.03	-0.18	-0.65	-0.69	-0.63
	Scenario 3	-1.15	-1.89	-0.23	-1.34	-1.95	-0.75	0.45	-0.80	6.69	0.12	-0.04	-0.35	-1.50	-1.58	-1.43
China	Scenario 1	-0.70	-1.05	-0.29	-0.64	-0.30	0.15	0.18	-0.44	3.13	0.18	-0.24	-0.06	-0.26	-0.21	-0.44
	Scenario 2	-0.68	-1.03	-0.29	-0.56	-0.30	0.15	0.00	-0.44	3.13	0.18	-0.25	-0.06	-0.26	-0.22	-0.44
	Scenario 3	-1.06	-1.56	-0.46	-0.87	-0.50	0.24	0.27	-0.66	4.85	0.28	-0.39	-0.08	-0.41	-0.33	-0.69
India	Scenario 1	-0.58	-0.90	-0.11	-0.36	-0.41	0.04	-0.13	-0.21	1.81	0.02	-0.13	-0.41	-0.27	-0.14	-0.24
	Scenario 2	-0.58	-0.90	-0.11	-0.36	-0.41	0.04	-0.13	-0.21	1.81	0.02	-0.13	-0.41	-0.27	-0.14	-0.24
	Scenario 3	-1.42	-2.17	-0.27	-0.90	-1.00	0.12	-0.26	-0.52	4.68	0.06	-0.28	-1.02	-0.61	-0.34	-1.05
Indonesia	Scenario 1	-0.13	-0.43	-0.19	-0.25	-0.01	-0.06	0.18	-0.07	1.92	0.03	0.01	-0.02	-0.12	-0.17	-0.27
	Scenario 2	-0.14	-0.44	-0.18	-0.25	-0.01	-0.06	0.18	-0.07	1.92	0.14	0.01	-0.02	-0.12	-0.18	-0.28
	Scenario 3	-0.34	-1.11	-0.47	-0.65	-0.02	-0.14	0.43	-0.18	4.77	0.35	0.05	-0.03	-0.31	-0.43	-0.68
Korea	Scenario 1	-0.47	-0.46	0.03	-0.57	-0.45	-0.11	-0.65	-1.00	4.19	0.15	-0.32	-1.08	-0.39	-0.42	-0.73
	Scenario 2	-0.57	-0.56	0.04	-0.70	-0.57	-0.14	-0.81	-1.23	5.16	0.18	-0.40	-1.35	-0.49	-0.52	-0.91
	Scenario 3	-0.71	-0.69	0.04	-0.89	-0.72	-0.18	-1.02	-1.53	6.59	0.25	-0.50	-1.68	-0.62	-0.65	-1.14
Vietnam	Scenario 1	-0.08	-0.28	-0.09	-0.20	0.02	-0.09	-0.39	-0.63	1.36	-0.03	-0.19	-0.22	-0.17	-0.18	-0.33
	Scenario 2	-0.08	-0.28	-0.09	-0.20	0.02	-0.09	-0.39	-0.63	1.36	-0.03	-0.19	-0.22	-0.17	-0.18	-0.33
	Scenario 3	-0.24	-0.79	-0.27	-0.54	0.09	-0.25	-1.07	-1.71	3.73	-0.09	-0.52	-0.55	-0.44	-0.49	-0.90

Source: Authors' calculations.

Table 14: Simulation Results and Percentage Changes in Sectoral Exports

		Grains & Crops	Meat and Lifestock	Extraction	Processed Food	Textiles	Manufacturing	Distribution	Utilities	Communication Services	Business Services	Financial Services	Water	Transport Services	Construction	Other Services
EU28	Scenario 1	0.51	0.31	0.37	0.16	0.68	0.11	0.41	0.33	-3.09	-0.21	-0.25	0.15	-0.03	0.19	0.25
	Scenario 2	0.71	0.46	0.55	0.26	1.03	0.13	0.51	0.48	-4.87	0.14	-0.39	0.19	-0.03	0.24	0.29
	Scenario 3	0.87	0.59	0.66	0.31	1.23	0.16	0.62	0.57	-5.88	0.17	-0.48	0.22	-0.05	0.29	0.35
Brazil	Scenario 1	0.25	0.43	0.08	0.23	0.52	0.13	-0.17	0.43	-1.53	-0.09	0.02	0.14	0.14	0.22	0.15
	Scenario 2	0.84	1.48	0.31	0.81	1.79	0.37	-0.56	1.41	-5.11	0.20	0.14	0.39	0.46	0.70	0.50
	Scenario 3	0.50	0.21	0.32	-0.03	-0.08	-0.21	-0.14	-0.09	-0.35	-0.14	-0.53	-0.35	-0.40	-1.07	-0.82
China	Scenario 1	1.30	2.45	0.82	0.95	0.54	-0.54	-1.07	0.67	-6.78	-0.89	-0.05	0.30	0.07	-0.35	0.13
	Scenario 2	1.29	2.39	0.81	0.81	0.53	-0.54	-0.61	0.67	-6.77	-0.89	-0.04	0.31	0.07	-0.34	0.13
	Scenario 3	2.00	3.63	1.32	1.26	0.88	-0.85	-1.65	0.99	-10.49	-1.37	-0.06	0.47	0.11	-0.54	0.20
India	Scenario 1	0.94	1.70	0.36	0.49	0.57	-0.35	0.33	0.33	-4.16	-0.34	0.00	0.33	0.12	0.10	0.46
	Scenario 2	0.94	1.70	0.36	0.49	0.57	-0.35	0.33	0.33	-4.16	-0.34	0.00	0.33	0.12	0.10	0.46
	Scenario 3	2.35	4.15	0.89	1.28	1.47	-0.87	0.72	0.85	-10.66	-0.82	-0.01	0.83	0.27	0.26	2.81
Indonesia	Scenario 1	0.46	0.80	0.22	0.33	0.50	-0.24	-0.30	-0.14	-4.63	-0.34	-0.18	-0.22	-0.10	-0.09	0.10
	Scenario 2	0.47	0.81	0.22	0.33	0.52	-0.24	-0.30	-0.13	-4.63	-0.80	-0.18	-0.21	-0.10	-0.09	0.11
	Scenario 3	1.20	2.08	0.56	0.90	1.32	-0.61	-0.71	-0.31	-11.50	-2.03	-0.50	-0.60	-0.21	-0.24	0.28
Korea	Scenario 1	1.08	1.15	1.14	0.68	1.26	0.07	-0.28	1.23	-7.43	-0.36	0.50	1.00	0.12	0.57	0.68
	Scenario 2	1.29	1.43	1.44	0.84	1.59	0.08	-0.34	1.51	-9.15	-0.43	0.64	1.28	0.16	0.71	0.87
	Scenario 3	1.59	1.71	1.83	1.06	2.03	0.10	-0.44	1.86	-11.70	-0.61	0.78	1.55	0.19	0.89	1.07
Vietnam	Scenario 1	0.14	0.51	0.11	0.22	0.10	-0.13	0.69	1.03	-3.77	-0.29	-0.09	0.62	-0.04	-0.03	0.30
	Scenario 2	0.14	0.51	0.11	0.22	0.10	-0.13	0.69	1.03	-3.77	-0.29	-0.09	0.62	-0.04	-0.03	0.30
	Scenario 3	0.43	1.46	0.33	0.58	0.37	-0.40	1.86	2.75	-10.37	-0.76	-0.30	1.67	-0.15	-0.09	0.79

Source: Authors' calculations.

Due to the comparative disadvantages that may arise from less innovation in these sectors as a result of tighter data regulations, countries with tighter data regulations are likely to become more import-dependent in the data-intensive services sectors over time. One should note that while changes in output and trade are rather low for other sectors, the general pattern of the results indicate a shift in production from the services and manufacturing to the primary sector as a result of restrictions on the flow of data.

Note that the CGE model does not account for dynamic effects, such as the impact of regulations on competition and innovative behaviour. The results of this analysis might therefore considerably underpredict the economic losses of regulations on the free flow of data and data localization. These losses would, for example, comprise efficiency losses resulting from reduced competition and economic inefficiencies due to greater distance of domestic

data services providers and data-intensive downstream providers to the global technology frontier.

DEVELOPMENTS AND FURTHER RESEARCH

RECENT DEVELOPMENTS IN DATA LOCALIZATION REGULATIONS

Since the study was conducted in 2013, the authors have carried out more extensive research covering more than 60 countries around the world to be publicly released in a database by mid-2016. It identifies several regulatory measures that include data localization requirements and recent trends in this policy field. For the countries studied in this chapter, a few legislative proposals were made and additional measures have been implemented. A detailed overview of relevant measures is provided in Table 15.

Table 15: Overview of Legislative Measures for the Countries under Study

Country	Law	Scope
Brazil	<ul style="list-style-type: none"> Law No 12.965 (Marco Civil), passed in March 2014 	<ul style="list-style-type: none"> The Brazilian government considered requiring Internet Service Providers to store information regarding Brazilian users only on local servers. The provision did not make it to the final version of Marco Civil.
China	<ul style="list-style-type: none"> Various laws and guidelines, including Guidelines for Personal Information Protection Within Public and Commercial Services Information Systems Standing Committee of the National People's Congress in China Decision on Strengthening Protection of Online Information Non-binding national standards related to personal information published by the Standardization Administration and the General Administration of Quality Supervision, Inspection, and Quarantine People's Bank of China Notice to Urge Banking Financial Institutions to Protect Personal Financial Information (Notice) China's Management Measures for Population Health Information 	<ul style="list-style-type: none"> A plethora of complex data privacy laws has made compliance very difficult for companies that collect personal information. Cross-border data transfer restrictions are imposed by various industry guidelines for the information-services sector. These guidelines may serve as a "regulatory baseline" for law enforcement authorities to assess whether or not a business is in compliance with Chinese data privacy laws. Banks and financial institutions are prohibited from storing, processing or analyzing any personal financial information outside China that has been collected in China. Population health information needs to be stored and processed within China. In addition, storage is not allowed overseas. Licensing system for online taxi companies that requires them to host user data on Chinese servers. Online maps are required to set up their server inside of the country and must acquire an official certificate.
India	<ul style="list-style-type: none"> Information Technology Act 2000 Information Technology (Reasonable Security Practices and Procedures and Sensitive Personal Data or Information) Rules of 2011 National Security Council Secretariat proposal for data localization of email services 	<ul style="list-style-type: none"> With its "Reasonable Security Practices and Procedures," the Indian government introduced a strict consent requirement that only allows for sensitive personal data to be transferred abroad that is necessary for the performance of a lawful contract between the body corporate (or any person acting on its behalf) and the provider of information or such transfer has been consented to by the provider of information. In February 2014, media reported on a leaked internal note from the National Security Council Secretariat, which shows that a three-pronged strategy with strong elements of data localization is being considered. The proposal included mandating all email providers to set up local servers for their India operations such that "all data generated from within India should be hosted in these India-based servers and this would make them subject to Indian laws" (Thomas 2014).

Country	Law	Scope
Indonesia	<ul style="list-style-type: none"> • Law No. 11 regarding Electronic Information and Transaction of 2008 • Government Regulation No. 82 of 2012 regarding the Provision of Electronic System and Transaction • Draft Regulation with Technical Guidelines for Data Centres • Circular Letter of Bank Indonesia No. 16/11/DKSP of 2014 regarding E-money Operations 	<ul style="list-style-type: none"> • Regulation 82 states that the storing of personal data and performing a transaction with the data of Indonesian nationals outside the Indonesian jurisdiction is restricted. This requirement would appear to apply particularly to personal data and transaction data of Indonesian nationals that is used within Indonesia and/or related to Indonesian nationals. • Regulation 82 requires “electronic systems operators for public service” to set up a data centre and disaster recovery centre in Indonesian territory for the purpose of law enforcement and data protection. • In the Annex of Circular Letter of Bank Indonesia No. 16/11/DKSP Year 2014 regarding E-money Operations, there is a requirement for all operators of e-money to localize data centres and data recovery centres within the territory of Indonesia.
Korea	<ul style="list-style-type: none"> • Regulations on Financial Institutions’ Outsourcing of Data Processing Business and IT Facilities’ approved in June 2013 • Spatial Data Industry Promotion Act 	<ul style="list-style-type: none"> • Despite provisions in its free trade agreements with EU and US to allow sending financial data across borders, Korea still prohibits outsourcing of data-processing activities to third parties in the financial services industry. Banks can therefore only process financial information related to Korean customers in-house, either in Korea or abroad, and offshore outsourcing is restricted to a financial firm’s head office, branch or affiliates. • Since June 2015, financial services institutions are allowed to offshore data processing to professional IT companies whose infrastructure is located outside of Korea. • Korea imposes a prohibition to store high-resolution imagery and related mapping data outside the country and justifies this restriction on security grounds.
Russia	<ul style="list-style-type: none"> • Federal Law No. 152-FZ “On Personal Data” as amended in July 2014 by Federal Law No. 242-FZ “On Amendments to Certain Legislative Acts of the Russian Federation for Clarification of Personal Data Processing in Information and Telecommunications Networks” • New provisions in the federal law on information, information technologies and protection of information (known as Blogger’s Law) • Federal Law No. 319-FZ “On Amendments to the Federal Law on the National Payment System and Certain Legislative Acts of the Russian Federation” 	<ul style="list-style-type: none"> • In accordance with the amendments to Federal Law No. 152-FZ of July 2006, an operator is required to ensure that the recording, systemization, accumulation, storage, clarification (updating, modification) and retrieval of Russian citizens’ personal data is to be conducted only in databases located within Russia. • The law affects all business practices that involve the processing of personal data of Russian citizens, irrespective of whether companies have a physical presence in Russia. • Blogger’s Law requires organizers of information distribution in the Internet (it is not clear which operators fall under this definition) to store on Russian territory information on facts of receiving, transfer, delivery and/or processing of voice information, texts, images, sounds and other electronic messages and information about users during 6 months from the end of these actions. • The amendments to the National Payment System Law require international payment cards to be processed locally.
Vietnam	<ul style="list-style-type: none"> • Decree No. 72/2013/ND-CP of July 15, 2013, on the Management, Provision and Use of Internet Services and Online Information 	<ul style="list-style-type: none"> • The Decree No. 72/2013 entered into force in September 2013 establishes local server requirements for online social networks, general information websites, mobile telecoms network-based content services and online games services. • All these organizations are required to establish at least one server inside the country “serving the inspection, storage, and provision of information at the request of competent state management agencies.”

Source: Authors.

For China, for example, a plethora of complex data privacy laws make compliance very difficult for companies that collect personal information. In addition, cross-border data transfer restrictions are imposed by various industry guidelines for the information services sector. These guidelines frequently serve as a “regulatory baseline” for law enforcement authorities to assess whether or not a business is in compliance with Chinese data privacy laws. Moreover, banks and financial institutions operating in China are prohibited from storing, processing or analyzing any personal financial information outside China that has been collected in China. The Vietnamese government imposed a decree establishing local server requirements for online social networks, general information websites, mobile telecoms, network-based content services and online games services. Affected organizations are required to establish at least one server inside the country “serving

the inspection, storage, and provision of information at the request of competent state management agencies.” (The Government of Vietnam 2013). As concerns Brazil, it is noteworthy that the Brazilian government considered forcing Internet Services Providers to store information regarding only Brazilian users on local servers, but respective provisions did not make it to the final version of the proposed law.

The landscape of legislative data localization requirements is highly diverse. Table 16 provides a preliminary excerpt of the ongoing research on data localization and affiliated measures beyond the set of countries studied in this chapter. Some countries are imposing local storage requirements — i.e., only a copy of the data has to remain within the territory of the country. This is the case in Denmark, Germany, Greece, the Netherlands, New Zealand, Poland, Romania, Russia, Sweden and Turkey. These measures

Table 16: Overview of Subjects Targeted by Data Localization Requirements (by country)

	Light (Only Copy)	Medium (Copy and Processing)	Strong (Ban to Transfer)
Australia			health data
Brunei		all data generated within the country	
Bulgaria		Gaming data	
Canada			data of public bodies
China		all data generated within the country, taxi users data, online maps, electronic media	financial information, health data, state secrets
Denmark	Financial records		
France		Systems for interception of electronic communication	
Germany	Tax records, accounting documents and business letters, invoices		
Greece	Data on ‘traffic and localisation’		
Indonesia		financial data	personal data
Korea			financial data, high resolution imagery and related mapping data
Luxembourg		Financial data	
Netherlands	Public records		
New Zealand	Business records		
Nigeria		Subscriber and consumer data, financial data	Government Data
Pakistan			Certain countries
Poland	Gambling data		
Romania	Gambling data		
Russia	Users information	Personal data	
Sweden	Certain corporate documents, certain public data		
Taiwan			China
Turkey	Online payments		
Vietnam		Online social networks, general information websites, mobile telecoms network based content services and online games services	

Source: Authors.

are usually imposed on a specific set of data relating to corporate documents, and the local storage is usually imposed so authorities can easily access such documents.

In other cases, countries are not only imposing local storage, but also local processing requirements. This represents a more restrictive barrier accompanied by high economic costs, as businesses are required to establish data servers in the implementing country or switch to local data services suppliers. Countries imposing such strict regimes include Brunei, China, France, Indonesia, Luxembourg, Nigeria, Russia and Vietnam. In most of these cases, the legislative measures cover a broad range of types of data. In the extreme case of Russia, they apply to all personal data — i.e., virtually all data transferred cross-border. The imposition of such a regime in Russia is quite recent. The Federal Law No. 152-FZ “On Personal Data” was, in fact, amended in July 2014 by Federal Law No. 242-FZ “On Amendments to Certain Legislative Acts of the Russian Federation for Clarification of Personal Data Processing in Information and Telecommunications Networks.” Such amendments, in force since September 2015, require data operators to ensure that the recording, systematization, accumulation, storage, update/amendment and retrieval of personal data of the citizens of the Russian Federation is made using databases located in the Russian Federation (Article 18 §5).

There are also few instances of bans to the free transfer of data. In these cases, companies not only have to store and process data within the border of the country, but they are also not allowed to send a copy of the data abroad. Such measures usually apply to especially sensitive data (as in the case of Australia where there is a ban to transfer health data abroad), but have also been used more extensively in two Canadian regions, China, Indonesia, Korea, Nigeria, Pakistan and Taiwan. For example, the two Canadian regions of British Columbia and Nova Scotia require that personal information held by a public body (primary and secondary school, universities, hospitals, government-owned utilities and public agencies) must be stored or accessed only in Canada. However, a public body may override the rules where storage or access outside of the respective province is essential. Moreover, the data can be transferred outside Canada “if the individual the information is about has identified the information and has consented, in the prescribed manner, to it being stored in or accessed from, as applicable, another jurisdiction.”¹⁴

Finally, it is important to note that conditional flow regimes — i.e., regimes under which certain conditions need to be fulfilled for data to leave the implementing

jurisdiction — can also effectively result in a ban to transfer data. These regimes can be so restrictive to cause a de facto ban to transfer specific data, as is the case in China. For personal data of European citizens, companies have the possibility to fulfill certain conditions required by legislators to transfer data abroad. Under the European Directive 95/46/EC, data is freely allowed to flow outside the European Economic Area only where:

- the recipient jurisdiction has an adequate level of data protection;
- the controller adduces adequate safeguards (for instance, by using model contract clauses, binding corporate rules or other contractual arrangements);
- the data subject has given his/her consent unambiguously;
- the transfer is necessary for the performance of a contract between the data subject and the controller;
- the transfer is necessary for the performance of a contract concluded in the interest of the data subject;
- (vi) the transfer is justified by public interest;
- the transfer is necessary to protect the vital interests of the data subject; and
- the data is public.

In addition to these options, the Safe Harbour agreement acted as a self-certification system open to certain US companies for the data protection compliance until its invalidation by the European Court of Justice in October 2015. Since then, there is a high level of legal uncertainty regarding data transfers to the United States. The European Commission has proposed a new regime (the so-called Privacy Shield) to replace the Safe Harbour. However, national Data Protection Authorities in the European Union have not yet expressed their opinion on the text, and the Article 29 Working Party will give a non-binding opinion on the privacy agreement and alternative mechanisms of data transfer only in mid-April 2016. Therefore, it remains a possibility that data transfer to the United States will be further restricted, as some Data Protection Authorities have already hinted at the possibility of imposing a ban to transfer data there.

SCOPE FOR FURTHER RESEARCH

This study is a first attempt to quantify the economic impact of several regulations of cross-border data flows and data localization measures. It applies an indirect top-down approach based on observable regulatory variables and econometric methods to calculate economic costs in terms of factor productivity losses. Using a CGE model, the losses in productivity have been translated into changes in aggregate economic

¹⁴ Freedom of Information and Protection of Privacy Act [RSBC 1996] CHAPTER 165, available at www.bclaws.ca/civix/document/LOC/complete/statreg/--%20F%20--/Freedom%20of%20Information%20and%20Protection%20of%20Privacy%20Act%20%5BRBC%201996%5D%20c.%20165/00_Act/96165_03.xml#section30.1.

activity, industrial output and industrial trade volumes. An indirect methodology, as applied here, is highly sensitive to the assumptions made for the degree of restrictiveness of the regulatory measures and the assumptions underlying the equational system of the applied CGE model.

This methodology is applied primarily due to the lack of sufficient data and sufficiently reliable information on the actual impact of certain data regulation policies at the business or industry level. Industry (survey) data for regulation-induced changes in the employment and cost of labour and capital — as well as data on the actual flows of data and data-intensive goods and services — would not only allow for greater precision in applying the indirect approach, it would also enable researchers to estimate sectoral and aggregate economic effects directly through the application of bottom-up instead of top-down approaches.

The methodology could also be improved by including different layers of data localization, from storage requirements to conditional flow regimes. In addition, a “right to be forgotten” legislation could be analyzed as a separate restriction. Again, reliable industry data on how these regulations affect businesses would significantly improve the empirical strategies applied.

CONCLUSION

The results demonstrate that communication services sectors show comparatively large productivity losses due to their high dependency on data inputs covered by data regulations. Data-intensive business and financial services also show relatively high losses in productivity. As concerns economic output, the production of data-intensive manufacturing and services sectors shrinks in all countries due to regulations on the free flow of data. Losses are notably taking place in the services sectors. The greatest declines in industry output are found for communications and business services, but also for financial services. At the same time, less data-intensive sectors are less affected by data regulations. The general patterns in the results indicate a shift in production from the services and manufacturing to the primary sector as a result of restrictions on the flow of data. Accordingly, tight regulations on the free flow of data tend to cause an economy’s production structure to shift (back) toward less innovative and relatively volatile sectors such as agriculture, raw materials and natural resources.

ACKNOWLEDGEMENTS

The authors would like to thank Sang-Seung Yi and Kyungsin Park during this research.

WORKS CITED

- Arnold, Jens Matthias, Beata Javorcik and Aaditya Mattoo. 2011. “Does Services Liberalization Benefit Manufacturing Firms? Evidence from the Czech Republic.” *Journal of International Economics* 85 (1): 136–46.
- Arnold, Jens Matthias, Beata Javorcik, Molly Liscomp and Aaditya Mattoo. 2012. “Services Reform and Manufacturing Performance: Evidence from India.” World Bank Policy Research Papers Series No. 5948. Washington, DC: World Bank.
- Bauer, Matthias, Hosuk Lee-Makiyama, Erik van der Marel and Bert Vershelde. 2014. “The Costs of Data Localization: Friendly Fire on Economic Recovery.” European Centre for International Political Economy Occasional Paper No. 3/2014, Brussels: ECIPE.
- Christensen, Laurits R., Andrea Colciago, Federico Etro and Greg Rafert. 2013. “The Impact of the Data Protection Regulation in the EU.” Intertic Policy Paper.
- European Commission. 2012. “Impact Assessment for the Regulation of the European Parliament and the Council on the Protection of Individuals with regard to the Processing of Personal Data and on the Free Movement of such Data.” Commission Staff Working Paper SEC(2012)72. Brussels: European Commission.
- Jorgenson, Dale W., Mun Ho and Jon D. Samuels. 2010. “Information Technology and US Productivity Growth: Evidence from a Prototype Industry Production Account.” In *Industrial Productivity in Europe: Growth and Crises*, edited by Matilda Mass and Robert Stehrer, 35–64. Northampton, MA: Edward Elgar Publishing.
- Jorgenson, Dale W., Mun Ho and Kevin J. Stiroh. 2005. *Information Technology and the American Growth Resurgence*. Cambridge, MA: MIT Press.
- . 2007. “Industry Origins of the American Productivity Resurgence.” *Economic Systems Research* 19 (3): 229–52.
- Koske, Isabell, Isabelle Wanner, Rosamaria Bitetti and Omar Barbiero. 2015. “The 2013 Update of the OECD Product Market Regulation Indicators: Policy Insights for OECD and non-OECD Countries.” OECD Economics Department Working Papers No. 1200.
- Le Merle, Matthew, Raju Sarma, Tashfeen Ahmed and Christopher Pencavel. 2012. *The Impact of EU Internet Privacy Regulations on Early Stage Investment: A Quantitative Study*. London: Booz & Co.
- Miroudot, Sébastien, Ranier Lanz and Alexandros Ragoussis. 2009. “Trade in Intermediate Goods and Services.” OECD Trade Policy Working Paper No. 93.

PriceWaterhouseCoopers. 2013. *Information Security Breaches Survey: Technical Report*. UK Department for Business Innovation and Skills. London: PWC.

Sáez, Sebastián, Daria Taglioni, Erik van der Marel, Claire H. Hollweg and Veronika Zavacka. 2015. "Valuing Services in Trade: A Toolkit for Competitiveness Diagnostics." Washington, DC: World Bank.

The Government of Vietnam. 2013. "Decree No. 72/2013/ND-CP of July 15, 2013, on the management, provision and use of Internet services and online information." www.vnnic.vn/sites/default/files/vanban/Decree%20No72-2013-ND-CP.PDF.

Thomas, Thomas K. 2014. "National Security Council proposes 3-pronged plan to protect Internet users." *The Hindu Business Line*, February 13. www.thehindubusinessline.com/info-tech/national-security-council-proposes-3pronged-plan-to-protect-internet-users/article5685794.ece.

UK Ministry of Justice. 2012. *Impact Assessment for the Proposal for an EU Data Protection Regulation*. London.

Van der Marel, Erik, Hosuk Lee-Makiyama, Mathhias Bauer and Bert Vershelde. forthcoming 2016. "A Methodology to Estimate the Costs of Data Regulation." *International Economics*. doi: 10.1016/j.inteco.2015.11.001.

ABOUT THE AUTHORS

Matthias Bauer is a senior economist at the European Centre for International Political Economy (ECIPE) in Brussels, Belgium. He studied business administration at the University of Hull, United Kingdom, and economics at the Friedrich Schiller University Jena, Germany. He received his Ph.D. after joining the Bundesbank graduate program on "Foundations of Global Financial Markets and Financial Stability."

Before joining the ECIPE, Matthias was the coordinator of international political economy at the international cooperation division of Konrad Adenauer Foundation, Berlin. He previously held positions at DekaBank, UBS and Mercedes-Benz China. He also works as a start-up and business development consultant.

Martina F. Ferracane is a policy analyst at ECIPE. Her work focuses on EU sectoral policies, especially in the areas of international trade, health care and digital innovation. She has worked in the area of private sector development at the European Commission, at the Trade and Investment Division of the UN Economic and Social Commission for Asia and the Pacific in Bangkok, and currently leads an association promoting digital fabrication in high schools and recently spent a few months in a FabLab in Brazil to learn digital fabrication techniques.

Martina holds a master's degree (honours) in economic internationalization, integration and international trade from the University of Valencia in Spain, and her bachelor's (honours) in economics and institutions of international and European integration from La Sapienza University of Rome.

Erik van der Marel is a senior economist at the ECIPE. His research is concentrated on empirical analyses in services trade. He is currently carrying out research in digital trade and the cross-border flows of data, as well as developing a database that covers the regulatory cost factors of digital trade, including data transfers across borders.

Prior to his appointment at ECIPE, Erik lectured at the London School of Economics, where he taught international political economy and the political economy of international trade at postgraduate level. He was also a research fellow at the Groupe d'Économie Mondiale in Paris at SciencesPo.

Erik received his Ph.D. in international economics from SciencesPo, where he specialized in the links between regulation and productivity and comparative advantage in services. He also holds an M.Sc. in economics from the Erasmus University in Rotterdam and an M.A. in European politics from the College of Europe in Bruges.

**SECTION THREE:
LEGAL JURISDICTION AND INTEROPERABILITY**

**CHAPTER FIVE:
JURISDICTION ON THE INTERNET:
FROM LEGAL ARMS RACE TO TRANSNATIONAL COOPERATION**

Bertrand de La Chapelle and Paul Fehlinger

Copyright © 2016 by Bertrand de La Chapelle and Paul Fehlinger

ACRONYMS

ccTLDs	country-code top-level domains
DNS	domain name system
G20	Group of Twenty
gLTDs	generic top-level-domains
ICANN	Internet Corporation for Assigned Names and Numbers
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISPs	Internet Service Providers
MLATs	mutual legal assistance treaties
OECD	Organisation for Co-operation and Development
UN	United Nations
UNESCO	United Nations Educational, Scientific and Cultural Organization
WSIS	World Summit on the Information Society

INTRODUCTION

In managing, promoting and protecting [the Internet's] presence in our lives, we need to be no less creative than those who invented it. Clearly, there is a need for governance, but that does not necessarily mean that it has to be done in the traditional way, for something that is so very different.

– Kofi Annan, then UN Secretary-General¹

The topic of jurisdiction has become a core issue for debate on the future of the Internet. The Internet's cross-border nature has produced unprecedented benefits for mankind. But it also generates tensions between national legal systems based on the territoriality of jurisdiction, particularly when dealing with abuses on the global network and Internet-related disputes.

Rooted in the treaties of the Peace of Westphalia of the seventeenth century, our international system is based on the separation of sovereignties, and these traditional modes of interstate cooperation struggle to cope with the digital realities of the twenty-first century.

We are confronted therefore with two major challenges: how to preserve the global nature of cyberspace while respecting national laws, and how to fight misuses and abuses of the Internet while ensuring the protection of

human rights. Both challenges require cooperation and clear procedures across borders to ensure efficiency and due process.

Since 2012, the Internet & Jurisdiction Project has provided a neutral dialogue space for a policy network comprising more than 100 key stakeholders from around the world to explore operational solutions for transnational cooperation on jurisdictional issues. This chapter directly draws upon the insights emerging from this pioneering multi-stakeholder process.

It addresses successively:

- why these issues represent a growing concern for all stakeholders, who are under pressure to find rapid solutions as the uses and misuses of the Internet increase;
- the legal arms race produced by the uncoordinated and unrestrained application of territoriality;
- the struggle of traditional modes of international cooperation to deal with this situation, especially with regard to access to user data, content takedowns and domain seizures;
- the resulting dangerous path that threatens to destroy the nature and benefits of the global network and the risks related to economy, human rights, infrastructure and security;
- the need to fill the institutional gap in Internet governance through innovative processes involving all stakeholder groups; and
- how to move toward transnational cooperation frameworks.

NATIONAL JURISDICTIONS AND CROSS-BORDER CYBERSPACES

CONFLICTING TERRITORIALITIES

The technical architecture of the Internet was conceived as cross-border and non-territorial from the onset. The World Wide Web technically allows, by default, access to any link regardless of physical location, and social media platforms serve hundreds of millions of users in shared cross-border online spaces. This transnational nature of the Internet has generated unprecedented benefits for humankind, be they political, economic or social. In particular, it uniquely fulfills the promises of Article 19 of the Universal Declaration of Human Rights regarding access to information “irrespective of frontiers.”

Yet, globally accessible content that is legal in one country may be illegal or even criminal in another. Like any human-made tool, the Internet is susceptible to misuse, and so,

¹ The UN Secretary-General's remarks at the opening session of the Global Forum on Internet Governance on March 24, 2004. www.un.org/sg/STATEMENTS/index.asp?nid=837.

cross-border cybercrime develops. Moreover, online communication tools are increasingly used by criminals “in the real world,” and access to information stored by Internet operators in other countries becomes essential in investigations.

From a historical perspective, cross-border interactions were rare, and international legal cooperation tools were designed to handle these exceptions. However, on the open Internet, interactions across borders are becoming the new normal. As a consequence, cross-border conflicts arise between users, the services they use, public authorities and any combination thereof. How to determine the applicable laws when interactions are transnational is becoming increasingly difficult, as the current international system is based on a patchwork of separate and territorially defined national jurisdictions.

Teresa Scassa and Robert J. Currie (2010) argue that, “put simply, because the Internet is borderless, states are faced with the need to regulate conduct or subject matter in contexts where the territorial nexus is only partial and in some cases uncertain. This immediately represents a challenge to the Westphalian model of exclusive territorial state sovereignty under international law.”

At least four territorial factors can play a role in determining applicable law: the location of the Internet user(s); the location of the servers that store the actual data; the locus of incorporation of the Internet companies that run the service(s) in question; and, potentially, the registrars or registries through which a domain name was registered.

These overlapping and often conflicting territorial criteria make both the application of laws in cyberspace and the resolution of Internet-related disputes difficult and inefficient. The principles of separation of sovereignties and non-interference between states that underpin the international system often render court decisions difficult to enforce and prevent the cooperation across borders necessary to efficiently deal with misuses online.

Tensions arise and will only grow as Internet penetration reaches four or five billion users from more than 190 different countries with diverse and potentially conflicting national laws and social, cultural or political sensitivities.

A CHALLENGE FOR ALL STAKEHOLDERS

The present situation is a concern for each category of actors.

Governments have a responsibility to ensure respect of the rule of law online, protect their citizens and combat crime. A sense of frustration prevails in the absence of clear standards on how to enforce national laws on the cross-border Internet. Law enforcement agencies in particular feel unable to conduct necessary investigations

to stop transnational crime and misuses of the network. In a system based on Westphalian territoriality, the principle of separation of jurisdictions becomes an obstacle to international cooperation.

Global Internet platforms, which relied on terms of service early on to establish the jurisdiction of their country of incorporation, now have to handle — and interpret — the 190-plus different national laws of the countries where they are accessible. This is a particular challenge to start-ups and medium-sized companies. Faced with growing direct requests for content takedown or access to user data, they also fear losing the protection of the limited-liability regime they have enjoyed so far and becoming responsible for thousands of micro-decisions of a quasi-judiciary nature² with significant human rights dimensions and reputation risks.

Technical operators worry that the fundamental separation of layers that forms the basis of the Internet architecture becomes blurred. Registries and registrars in particular see increasing efforts to leverage the domain name system (DNS) as a content control tool with global reach. Hosting providers and internet service providers (ISPs) are equally concerned.

Civil society groups around the world worry about a potential race to the bottom in terms of protection of freedom of expression and privacy and a perceived privatization of dispute resolution. **Average users** are confused by the legal uncertainty about what rules apply to their online activities and feel powerless to obtain predictable and affordable redress when harmed, as multi-national litigation is beyond their reach.

International organizations struggle because of overlapping thematic scopes, or a geographical remit that is not universal. Although some, such as the Council of Europe, the Organisation for Economic Co-operation and Development (OECD), and the United Nations Educational, Scientific and Cultural Organization (UNESCO) have made significant efforts to include civil society, the private sector and the technical community in their processes, they remain by nature intergovernmental organizations. As such, they are limited in their capacity to put sensitive but necessary issues on their agenda by the lack of consensus, or worse, dissent among their members.

A CORE ISSUE OF INTERNET GOVERNANCE

The jurisdictional challenge is at the nexus of Internet governance and touches upon multiple traditional policy areas: the development of the global digital economy, ensuring a clear and predictable legal environment through cooperation, guaranteeing the exercise of fundamental

² Jacques de Werra (2015) labelled this new phenomenon “massive online micro justice.”

human rights, and ensuring security and public order. Since 2012, the Internet & Jurisdiction Project's Observatory has documented more than 1,000 high-level cases around the world that show the growing tension between national jurisdictions³ due to the cross-border nature of the Internet.

Contrary to what they may perceive, however, the different categories of stakeholders have less of a problem with each other than a problem in common — that is, how to manage the coexistence of different norms in shared online spaces. Realizing this is the necessary first step toward a common solution. As the World Economic Forum's 2016 report on Internet fragmentation shows, trends toward the re-nationalization of cyberspaces are observable (Drake, Cerf and Kleinwächter 2016). Maintaining a global Internet *by default*, which fulfills the ambitions of the Universal Declaration of Human Rights, notably article 19, and boosts innovation and growth through online services and the cloud economy, requires transnational legal cooperation.

Within the global Internet & Jurisdiction multi-stakeholder process, three key issues have emerged as potential areas for such cooperation:

- **Domain name seizures:** Under which conditions and criteria is action at the DNS level justified, given its global impact?
- **Content takedown and withholding:** How can stakeholders determine proportionate restrictions to access that respect both national laws and international human rights?
- **Access to user data:** Under which conditions can law enforcement in one country obtain communication of user information from a foreign operator?

In each case, both procedural and substantive elements need to be addressed to develop balanced regimes.

Unfortunately, unilateral actions by actors to solve the complex jurisdictional conundrum on their own create a legal competition that makes the problem harder, rather than easier, to solve.

A LEGAL ARMS RACE IN CYBERSPACE?

Solving the Internet and jurisdiction challenge is intrinsically linked to the general debate about modalities of global governance. Christoph Knill and Dirk Lehmkuhl (2002) already observed in 2002 that “[e]conomic and technological interdependencies have created a range of problems that exceed the scope of national sovereignty and can therefore no longer be sufficiently resolved by the unilateral action of national governments.”

3 See the Internet & Jurisdiction Observatory Retrospect Archive (n.d.).

Yet, confronted with increasing domestic pressure to address cyber issues, governments feel compelled to act on their own, using an extensive interpretation of territoriality criteria. This “hyper-territoriality” manifests itself by either extending sovereignty beyond national frontiers or reimposing national borders.

EXTRATERRITORIALITY

Extraterritorial extension of national jurisdiction is becoming the *realpolitik* of Internet regulation.

First of all, governments with Internet platforms or technical operators incorporated on their soil can impose their national laws and regulations on these private actors, with direct transboundary impacts on all foreign users of these services. An often-cited example regarding the United States is the surveillance capacities described in the Snowden revelations. Regarding the reach of law enforcement, an ongoing landmark lawsuit will determine whether US authorities have a right to access emails stored by Microsoft, a US company, in its data centres in the Irish jurisdiction.⁴ Previous cases involved a Department of US Homeland Security agency seizing domain names belonging to foreign registrants on the sole basis of their registration through a US-based registrar (the RojaDirecta case⁵) or registry (the Bodog case⁶).

Furthermore, draft legislations increasingly include clauses establishing extraterritorial reach, such as the UK Investigatory Powers Bill⁷ or the General Data Protection Regulation in the European Union.⁸

Finally, litigation also plays a prominent role in setting new global standards, with impacts far beyond the respective jurisdictions. Facebook, for instance, changed its global terms of service after a US court decision on its “sponsored stories” feature.⁹ Courts increasingly affirm competence regarding services incorporated in other countries merely because they are accessible in their territory, as illustrated by the recent Yahoo case in Belgium.¹⁰ Some difficulties naturally exist in enforcing the resulting judgments, as the national blockade of WhatsApp in Brazil showed.¹¹ Yet local cases can have global impacts. For instance, after the Court of Justice of the European Union Costeja decision

4 See Internet & Jurisdiction Retrospect (2015).

5 See Internet & Jurisdiction Retrospect (2012a).

6 See Internet & Jurisdiction Retrospect (2012b).

7 See Internet & Jurisdiction Retrospect (2016).

8 See Internet & Jurisdiction Retrospect (2015b).

9 See Internet & Jurisdiction Retrospect (2013).

10 See Internet & Jurisdiction Retrospect (2015c).

11 See Internet & Jurisdiction Retrospect (2015d).

(the right to be de-indexed), the French data protection authority demanded that Google extend its de-indexing to all versions of its search engine, arguing that the service is based on a single processing of data worldwide.¹²

Local court decisions can also trigger new international norms for the interaction between states and Internet companies. For instance, the right to be de-indexed, initially established by Europe for Google, is now implemented by other search engines such as Microsoft Bing or Yahoo Search¹³ and has produced ripple effects in Asia¹⁴ and Latin America.¹⁵

DIGITAL SOVEREIGNTY

Not all countries are able — or trying — to extend their sovereignty beyond their borders. As a consequence, re-nationalization is a complementary trend to extraterritorial extension of sovereignty. The theme of “digital sovereignty” gains traction in many jurisdictions in a context of rising tensions and a sense of powerlessness by public authorities to impose respect for their national laws on foreign-based Internet platforms and technical operators. This can mean efforts to literally re-erect borders on the Internet through blocking of uniform resource locators or Internet Protocol (IP) addresses via national ISPs — something that has become much easier to implement today than in the early 2000s — or the creation of a limited number of national gateways.

So-called “data localization” laws are also part of this trend. They range from indirect requirements that would impose data localization only as a last resort if companies fail to honour legitimate national requests (see Brazil’s Marco Civil¹⁶) to strict requirements, which stipulate that the data of national citizens processed by foreign companies needs to be stored within the national jurisdiction (see Russia¹⁷).

Other digital sovereignty measures can range from strong national intermediary liability regimes,¹⁸ requirements to open local offices, demanding back doors to encryption technologies or the imposition of full-fledged licensing regimes.

PARADOXES OF SOVEREIGNTY

Extreme and unrestrained leveraging of traditional territorial criteria introduces two paradoxes.

First, as described above, national actions upon operators with global reach have impacts on other jurisdictions. Such actions appear contrary to the very principle of non-interference, which is a direct corollary of sovereignty itself. This increases interstate tensions and potential conflicts between jurisdictions. While rewarding the most powerful digital countries, it encourages others to react and adopt measures based on mistrust and the reimposition of national borders.

Second, strict digital sovereignty measures such as data localization are not scalable globally. It is highly unlikely that necessary data centres could be, for example, established in all developing or small countries. Furthermore, although often presented as a tool to prevent surveillance, it might increase the likelihood of surveillance through the replication of data, which is required to create local copies that are stored in the reach of national authorities, while still allowing global processing and cross-border interactions.

Sovereignty is relevant in the digital age, but it behooves governments to take into account the potential transborder impact of their national decisions. This is why the recommendation adopted in 2011 by the Committee of Ministers of the Council of Europe established the responsibility of states to avoid “adverse transboundary impact on access to and use of the Internet” when they enforce national jurisdictions (Council of Europe 2011).

Exercised without restraint, both “extraterritorial extension of sovereignty” and “digital sovereignty” measures run contrary to the Kantian categorical imperative that should underpin international Internet regulation: Any national policy measure that would be detrimental if generalized around the world should not be adopted in the first place. International norms of cooperation are needed to prevent this legal arms race.

LIMITS TO INTERNATIONAL COOPERATION

Managing cross-border commons poses systemic difficulties for the existing international system (Ostrom 1990). The Westphalian principles of separation of sovereignties and non-interference actually represent more of an obstacle than a solution for cooperation on cyber issues.

John Palfry and Urs Gasser et al. (2012) and Rolf H. Weber (2014) rightfully argue that we need more legal interoperability to preserve the global nature of the Internet, but substantive harmonization of laws related to the use of

12 See Internet & Jurisdiction Retrospect (2015e).

13 See Internet & Jurisdiction Retrospect (2015f).

14 See Internet & Jurisdiction Retrospect (2014a).

15 See Internet & Jurisdiction Retrospect (2015g).

16 See Internet & Jurisdiction Retrospect (2014b).

17 See Internet & Jurisdiction Retrospect (2015h).

18 For an overview of national intermediary liability regimes, see Stanford World Intermediary Liability Map at <https://cyberlaw.stanford.edu/our-work/projects/world-intermediary-liability-map-wilmap>.

the Internet seems unattainable. Multilateral efforts have proved so far inconclusive; bilateral arrangements such as mutual legal assistance treaties (MLATs) are in dire need of reform; and the increasing number of informal interactions between public and private actors across borders lack procedural guarantees.

OBSTACLES TO MULTILATERAL EFFORTS

The Internet is by nature disruptive, including with respect to the international regulatory system. As A. Claire Cutler (2001, 133) puts it, “traditional Westphalian-inspired assumptions about power and authority are incapable of providing contemporary understanding, producing a growing disjunction between the theory and the practice of the global system.”

The idea of a global, all-encompassing Internet treaty that would harmonize relevant laws and solve the full range of cyber-cooperation issues is advocated only by some rare actors, who have tried to draw an analogy to decades-long efforts of international negotiations that resulted in the Law of the Sea Convention or the Outer Space Treaty. But the Internet is not a natural commons and, as Wolfgang Kleinwächter (2001) has argued, “while all these international conventions can be seen as great achievements of contemporary international law, it is hard to believe that this is a usable model for policy and law-making for the global Internet” due to the newness, volatility and rapid pace of innovation in the digital realm (Nye 2014).

Since the end of the World Summit on the Information Society (WSIS), intergovernmental discussions in various UN fora have made little progress beyond the wording of the Declaration adopted in Tunis in 2005. Moreover, the international community was split in 2012 during the World Conference on International Telecommunications, signifying the absence of global consensus not only at the level of substance, but even on the proper institutional framework for such discussions.

In any case, treaty negotiations are notoriously long. Even the most extensive agreement to date tackling cybercrime, the Budapest Convention, was a lengthy process. If formal negotiations took only four years, more than a decade was necessary to actually put the topic on the agenda. Although now signed by more than 50 states around the world (excluding, however, several large countries such as Brazil and India), some countries use the fact that it was elaborated initially within the Council of Europe as an argument to refuse joining a regime they did not participate in drafting. The Budapest Convention also require signatories to transpose its provisions into national laws and its Article 18 on “subscriber information” or Article 32b addressing “trans-border access to stored data” are often considered not sufficient enough to provide effective cooperation. Like all international agreements,

the Budapest Convention is also difficult to modify in response to rapidly changing technology.

In the past few years, many useful declarations have been developed within multilateral organizations at the level of general principles, showing some form of convergence. Still, none of them were able to move toward developing an operationally implementable regime.

MLATS: THE SWITCHED NETWORK OF INTERNATIONAL COOPERATION

Historically, the so-called MLATs enabling government-to-government legal cooperation were negotiated to handle rare and rather exceptional cross-border criminal cases. These intergovernmental tools allow public authorities in country A to ask for assistance to, for instance, access user data stored by an operator in country B. Upon receipt of the request, country B examines if it is also valid according to its national laws. If so, the data holder in country B is lawfully compelled to submit the data to authorities in country B, which will then share it with the requesting authorities of country A.

However, now that cross-border is the new normal on the Internet, this system is generally described as “broken.” MLATs have at least four structural limitations:

- **Speed:** MLATs are ill adapted to the speed of the Internet and the viral spread of information. In the best cases, an MLAT request from one government to another takes months to be processed. It can take up to two years between certain countries. The very elaborate circuit of validations is legitimately intended to provide procedural guarantees, but makes the whole system impracticable.
- **Scope:** MLATs are often limited to “dual incrimination” cases, that is, they cover only issues qualified as a crime in the jurisdictions of both requesting and receiving countries. Given the disparity of national legislations, their relevance is limited, particularly on speech issues (such as hate speech and defamation). They are also ineffective when the location of the data is unknown.
- **Asymmetry:** Regardless of the actual physical location of events or involved parties, the MLAT system de facto imposes the law of the recipient country over the law of the requesting one, even if there is no other territorial connection to the latter than the incorporation of the targeted platform or operator. An increasing number of countries find this unbalanced, given the dominant role of US-based companies.
- **Scalability:** The system of traditional MLAT treaties can hardly encompass the scale of the Internet. A large number of countries around the world do not have MLAT treaties with each other, and establishing such

bilateral relations among 190 countries would require more than 15,000 arrangements.¹⁹

The MLAT system is the switched network of international cooperation.²⁰ It is in dire need of reform to adapt to the Internet age and reforming it will not be easy. It will require more than simply streamlining existing procedures: creative solutions are needed to address its structural limitations and ensure both transnational due process and efficiency.

Recent initiatives have been launched in the United States, in particular to address the asymmetry issue, including a potential reform of the Electronic Communications Privacy Act of 1986. This represents a positive signal and international discussions are ongoing. The question of scope, however, remains, and many issues cannot be addressed via the MLAT approach as long as national legislations remain unharmonized.

THE RISE OF DIRECT PUBLIC-PRIVATE REQUESTS ACROSS BORDERS

In the absence of appropriate international cooperation frameworks, there are an increasing number of requests that public authorities in one country directly send to private actors in other jurisdictions, for the following three actions:

- **Domain seizures:** Removal of the entire domain of an allegedly infringing website.
- **Content takedown:** Removal or withholding of a specific piece of infringing content.
- **User data access:** Access to user information related to who posted infringing content, or other investigations.

There is a lack of reliable data to show the entire magnitude of this new trend. Transparency reports of some major global Internet companies provide a snapshot of the rise of such requests, but without sufficient harmonization of reporting methodologies. So far, only a small number of — mostly US-based — Internet companies publish such reports. Aggregated data from states, that is, the senders of these requests, is still unavailable. It is also important to understand that the original sending countries of MLAT requests are not revealed in such transparency reports, as these requests are ultimately handed down to companies as national requests from their respective countries of incorporation.

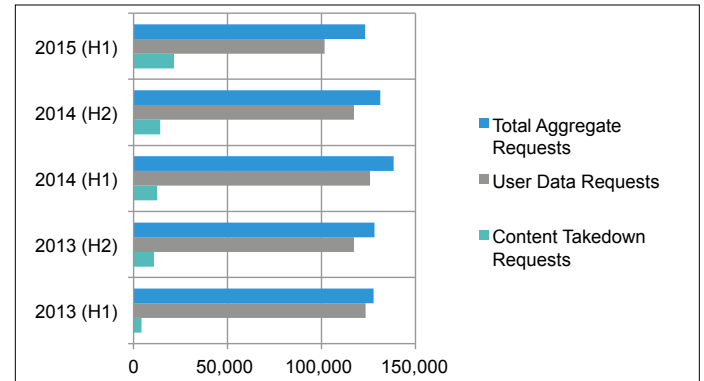
Pioneered by Google in 2009, transparency reporting is still a nascent trend. For example, nine out of the 13 analyzed

¹⁹ For an overview of existing MLAT treaties, consult the MLAT Map by the non-governmental organization Access Now, available at <https://mlat.info/>.

²⁰ For a comparison between the public switched telephone network and the distributed architecture of Internet routing see Internet Society (n.d.).

platforms only launched transparency reports in 2013. Nevertheless, Figure 1 provides an indicative statistical overview by showing a survey of the combined number of requests received from public authorities (courts, law enforcement, other agencies) as reported by 13 Internet platforms²¹ for content takedown and user data between 2013 and mid-2015.

Figure 1: The Rise of Direct Requests



Data sources: See footnote 21.

Note: H1 = first half of year, H2 = second half of year

Since 2013 the surveyed platforms reported in total 648,544 content removal requests (excluding copyright-related requests) and user information requests. The vast majority of reported requests have been addressed to four companies: Facebook, Google, Microsoft and Yahoo. The actual volume of such requests around the world is estimated to be much higher and will certainly rise with the next billion Internet users from increasingly diverse jurisdictions as they start using numerous Internet platforms and services.

Just in the first six months of 2015, Facebook (2015), for example, received requests from courts, law enforcement or other authorities from 92 jurisdictions, Google (2015) from 91 jurisdictions, Microsoft (2015) from 64 jurisdictions, Twitter (2015) from 37 jurisdictions and Yahoo (2015) from 34 jurisdictions.

This trend reflects an effort to establish modalities of voluntary cooperation between public and private actors across borders. However, it forces private entities to make determinations on sensitive high-stake issues regarding freedom of expression, human rights, economic conduct, international diplomacy and public safety

²¹ Combined data from transparency reports between 2013 and the first semester of 2015 on content takedown request (excluding copyright) and user information requests issued by governments (law enforcement, courts, other authorities) as reported by AOL (transparency reporting since 2011), Apple (since 2013), WordPress (since 2013), Dropbox (since 2013), Facebook (since 2013), Google (since 2010, although reports started in 2009), LinkedIn (since 2011), Microsoft (since 2013), Pinterest (since 2013), Snapchat (since 2014), Tumblr (since 2013), Twitter (since 2012), Wikimedia (since 2012) and Yahoo (since 2013).

through procedures and criteria that lack transparency and due process. It also often places them in a difficult situation, as when accepting a request conflicts with the law of their country of incorporation (for instance, direct communication of user content is prohibited by the Electronic Communications Privacy Act in the United States). Meanwhile, requests not honoured can lead to tensions or, in extreme cases, to the blocking of entire platforms by national ISPs or forced data localization. While world-leading platforms can afford to allocate the necessary human and financial resources, start-ups and medium-sized companies with globally available content and services have a greater struggle with this situation.

A DANGEROUS PATH

The lack of coordination and the inability of the Westphalian international system to provide the necessary cooperation solutions produce a typical “prisoner’s dilemma” situation. That is, every single actor, forced to use the only tools available to it, makes short-term decisions that appear in its immediate interest, though their cumulative effect is at best suboptimal and most likely detrimental to all in the longer term.

If we continue to lack appropriate cooperation mechanisms and “fall back into managing the national, rather than managing shared cross-border online spaces in a collaborative way” (Fehlinger 2014), the sum of uncoordinated unilateral actions by governments and private actors can have unintended consequences, with strong negative impacts in economic, human rights, infrastructure and security areas.

Unintended Consequences

ECONOMY	HUMAN RIGHTS
Demise of globally accessible services	Reduced freedom of expression across borders
Market entry barriers	Limits to access to information
Reduced investment in start-ups	Limits to freedom of assembly in cross-border online spaces
Stifled innovation	Lack of access to justice and redress
Disadvantages for developing countries	
INFRASTRUCTURE	SECURITY
Blurred separation of layers	Eroding of global cyber security
Facilitation of surveillance	Diplomatic tensions
Encryption wars	Increase of cybercrimes and online terrorism
Restrictions to the use of virtual private networks	Threats to human security
Reduced network resilience	

Source: Author.

ECONOMIC IMPACTS

In 2014, the Boston Consulting Group estimated the value of the digital economy of the Group of Twenty countries alone at US\$4.2 trillion, representing 5 to 9 percent of total GDP in developed countries (Boston Consulting Group 2014). The cross-border nature of the Internet and its cloud-based services are at the heart of innovation and growth. This is why the OECD is addressing the challenges to Internet openness in its June 2016 Ministerial Conference in Mexico and why the 2016 World Economic Forum’s Davos meeting discussed the impact of cyberspace fragmentation. A legal arms race and lack of cooperation would stifle innovation and competition, and jeopardize growth. Most established Internet companies were able to scale up internationally before the current move toward re-territorialization. The future development of global services and the cloud approach are at stake.

Investment in start-ups and medium-sized companies (especially those dealing with user-generated content) would decrease because of higher intermediary liability risks and legal uncertainty. Compulsory data localization might constitute a potential market entry barrier. Such requirements could be respected only by large, already established operators, limiting innovation and market accessibility for small companies wanting to serve a global market, particularly from developing countries.

HUMAN RIGHTS IMPACTS

International organizations such as UNESCO (“Internet universality”) or the Council of Europe (“cross-border flow of Internet traffic and Internet freedom”) have established the connection between human rights and the cross-border Internet (UNESCO 2013; Council of Europe 2015). It has uniquely fulfilled the promises of Article 19 of the Universal Declaration of Human Rights, allowing everyone to “seek, receive and impart information and ideas through any media and regardless of frontiers” (UN Human Rights Office of the High Commissioner (2011). enriched the social fabric across borders and improved our quality of life. Personal communication capacities are augmented, allowing frictionless expression, deliberation, and the holding of opinions across borders. The cross-border Internet facilitates the sharing and pooling of resources, and provides diasporas with irreplaceable communication tools. It has enabled the creation of critical-mass communities with common interests for social, political, or economic issues regardless of spatial distance and facilitated collaborative not-for-profit activities that have created tremendous global social value, such as Wikipedia.

The uncontrolled reterritorialization of the Internet in order to address its misuses could destroy the unprecedented human rights benefits the Internet has generated. Ironically, measures such as data localization and decryption could

in fact increase opportunities for surveillance rather than reduce them, as well as harm the right to privacy (UN Human Rights Office of the High Commissioner 2015). Increased pressure on Internet companies to accept direct requests could produce a “race to the bottom” by limiting freedom of expression and lowering due process protections. Conversely, the continued absence of affordable cross-border appeal and redress mechanisms for harmed Internet users has a serious negative impact on global justice.

TECHNICAL INFRASTRUCTURE IMPACTS

In 2013, the leaders of the 10 organizations responsible for coordination of the Internet’s technical infrastructure met in Montevideo, Uruguay, to stress in their joint statement “the importance of globally coherent Internet operations, and warned against Internet fragmentation at a national level”(Internet Corporation for Assigned Names and Numbers [ICANN] 2013). In enforcing national laws online in the absence of international cooperation frameworks, there is a temptation to use the technical infrastructure of the Internet to address content issues. This, however, blurs a fundamental architectural principle of the Internet: the separation of the neutral logical layer (DNS, IP addresses, et cetera) and the application layer (online platforms and services).

Leveraging the location of registries and registrars to impose the national laws of their country of incorporation on the global content under the country-code top-level domains (ccTLDs) or generic top-level-domains (gTLDs) they manage would be a clear extraterritorial extension of sovereignty, given the global impact of a domain seizure. In parallel, generalizing geo-IP filtering to withhold content on specific territories may lead to forcing Regional Internet Registries to systematically allocate IP addresses on a territorial basis. Such a scenario could complicate routing. With the transition from IP version 4 (IPv4) to IP version 6 (IPv6), it could even facilitate surveillance, should IP addresses be permanently hardwired to specific devices and become identity identifiers.

In an effort by Internet companies to reduce their multi-jurisdictional liability, unbreakable encryption technologies might lead to a spiral of encryption/decryption conflicts between public and private actors. The imposition of a limited number of Internet gateways to connect a territory in order to facilitate blocking measures potentially reduces the resilience of the overall technical network. Finally, the banning of technologies such as virtual private networks is not only contrary to Article 13(2) of the Universal Declaration of Human Rights,²² it also reduces the security of transactions and communications.

²² Universal Declaration of Human Rights Article 13(2): “Everyone has the right to leave any country, including his own, and to return to his country.”

SECURITY IMPACTS

The absence of agreed-upon frameworks to handle requests across borders has already resulted in diplomatic tensions between a country seeking to enforce its national laws and the country in whose jurisdiction the Internet platform or technical operator is actually located. Examples are Google’s China exit in 2010 (McCullagh 2010), the Indian Assam riots in 2012,²³ the Innocence of Muslim YouTube video in 2012²⁴ and Turkey’s blocking of Twitter in 2014.²⁵ Likewise, debates about MLAT reform are fuelling interstate dissonances. Such international conflicts are likely to increase if nothing is done.

It is the duty of states to protect their citizens and maintain public order within the provisions of Article 29 of the Universal Declaration of Human Rights. However, the rapid and viral propagation of incitation to violence (often called “digital wildfires”) could lead to disaster if we lack efficient transnational cooperation mechanisms that set standards and procedures for the interactions between states, Internet platforms and users across borders in situations of public order tensions. The international fight against terrorism online is emblematic of this challenge. Meanwhile, cybercrime is on the rise, and most online crimes have a multi-jurisdictional footprint, which makes cooperation across borders necessary to guarantee the security online, as well as off-line. The absence of appropriate regimes to access data across borders further increases the incentives for direct surveillance. Failure to develop the needed frameworks might ultimately lead to a decrease in global cyber security and order.

FILLING THE INSTITUTIONAL GAP IN INTERNET GOVERNANCE

Traditional intergovernmental cooperation mechanisms are failing so far to provide appropriate solutions. Legal harmonization on substance is difficult to achieve but the costs of inaction are daunting. There is an institutional gap in the Internet governance ecosystem that must be filled to adequately address these new challenges. In doing so, following the words of former UN Secretary-General Kofi Annan, we need to be as creative as the inventors of the Internet. To preserve the global nature of the Internet and address its misuses demands the development of innovative cooperation mechanisms that are as transnational, inclusive and distributed as the network itself.

²³ See Internet & Jurisdiction Retrospect (2012c).

²⁴ See Internet & Jurisdiction Retrospect (2012d).

²⁵ See Internet & Jurisdiction Retrospect (2014c).

LESSONS FROM THE TECHNICAL GOVERNANCE “OF” THE INTERNET

Internet governance was famously defined in the United Nation’s WSIS Tunis Agenda (2005) as “the development and application by governments, the private sector and civil society, in their respective roles, of shared principles, norms, rules, decision-making procedures, and programmes that shape the evolution and use of the Internet.”

In this definition, we see a distinction between governance “of” the Internet and governance “on” the Internet (de La Chapelle 2007). Governance “of” the Internet designates the governance of protocols, standards, addresses and *the evolution* of the technical architecture. Governance “on” the Internet relates to *the use* of the Internet, that is, the applications and services that run on top of the physical and logical layers, as well as Internet users’ behaviour. The jurisdictional challenges discussed in this chapter are primarily related to governance “on” the Internet.

A complex and robust network of institutions has emerged over time to handle governance “of” the Internet. It comprises, *inter alia*, the Internet Engineering Task Force and World Wide Web Consortium (W3C) for the development of Internet and web standards; five Regional Internet Registries allocating IP addresses; the 13 root servers and their multiple mirrors; ICANN; and the numerous registries and registrars distributing second-level domain names.

In dealing with the Internet’s logical layer, each of these institutions covers the five stages necessary for the “development and application” of governance regimes: issue-framing, drafting, validation, implementation and reviews. The policies developed through their bottom-up participatory processes can have wide-ranging transnational implications, such as when ICANN regulates the allocation of the semantic spectrum of gTLD extensions or the accreditation of market operators (registrars and registries).

Together, these institutions formed the necessary ecosystem of governance that has enabled the Internet to grow from the limited ambit of its research background to serve several billion people and permeate almost all human activities. This ecosystem of transnational institutions is fundamentally distributed; each entity deals with a specific issue, with loosely coupled coordination. It was developed progressively through time as policy needs arose. Each entity has its own specific institutional structure and internal procedures. Most important, they operate on a fundamental principle: the open participation of all relevant stakeholders in the processes dealing with issues they impact or are impacted by.

EVOLUTION OF THE ECOSYSTEM: GOVERNANCE “ON” THE INTERNET

By contrast, the institutional ecosystem addressing issues related to governance “on” the Internet is embryonic at best, or as Mark Raymond and Laura DeNardis (2015) elegantly expressed, “inchoate.”

The IGF is the main outcome of the WSIS process. In its 10 years of existence, it has demonstrated its capacity to act every year as a “watering hole,” where all actors identify challenges, share experiences and present their work. However, despite its undeniable success and essential role, not to mention the emergence of numerous national and regional spinoffs, it still only covers at best the first stages of the policy-making cycle: agenda setting and issue framing. Beyond some noteworthy efforts to document best practices, no efficient mechanisms exist yet to enable ongoing intersessional work on specific issues to produce, let alone implement and enforce, the needed transnational arrangements for governance “on” the Internet.

The NETmundial Roadmap, an outcome of the major 2014 multi-stakeholder conference, highlighted the jurisdiction issue as an important topic for the global community (NETmundial 2014). To preserve the cross-border nature of the Internet by default for the next generations to come, we need to collectively fill the institutional gap for the governance “on” the Internet. This is in line with the ambitions of the global Internet governance community to “further develop the Internet governance ecosystem to produce operational solutions for current and future Internet issues,” and to preserve the Internet as a “unified and unfragmented space” in a collaborative manner (NETmundial n.d.).

In doing so, we need to keep in mind the lessons that made the success of the existing institutional ecosystem for governance “of” the Internet. The robustness of the policies and solutions it produces is directly related to its fundamental characteristic of being transnational, open and organized in a distributed way. Given the diversity of the modes of organization of technical governance organizations, this does not mean the mere replication of a single model, but rather taking adequate inspiration from these principles to develop the governance “on” the Internet.

In the specific case of developing new transnational cooperation mechanisms for domain seizures, content takedowns and access to user data, the institutional gap of governance “on” the Internet lies at the intersection of four policy areas: legal interoperability, economy, human rights and cyber security (See Figure 2).

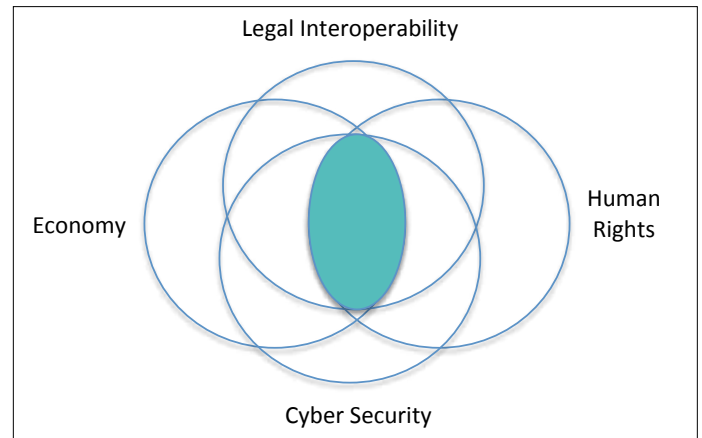
ENABLING ISSUE-BASED MULTI-STAKEHOLDER COOPERATION

The multi-stakeholder approach was explicitly endorsed by more than 180 countries at the heads of state level in the Tunis Agenda in 2005, and reconfirmed in the United Nations General Assembly High-Level Meeting on the WSIS+10 in December 2015. Filling the institutional gap requires neither the creation of new international organizations nor giving a unique responsibility to any existing one, as Internet issues are relevant to the mandates of a plurality of entities. A more creative approach is needed: the formation of issue-based governance networks.

In line with the 2014 recommendations of the High-Level Panel on Global Internet Cooperation and Governance Mechanisms (ICANN 2014) chaired by the President of Estonia, Toomas Ilves, developing transnational mechanisms for policy cooperation requires ongoing, multi-stakeholder and issue-based processes:

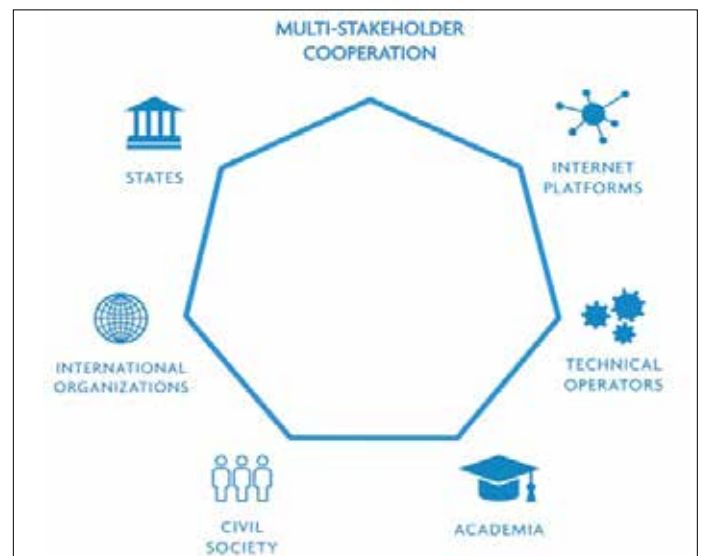
- **Ongoing**, because the current proliferation of one-shot conferences, fora, panels and workshops, however useful to foster mutual understanding, is not sufficient to move toward operational solutions. Developing networks, trust and a common approach to issues and objectives cannot be achieved in disconnected series of two-hour sessions.
- **Multi-stakeholder**, because no single stakeholder group working alone can grasp all the technical, political, legal, security, social and economic dimensions of an issue — a condition for the development of balanced regimes. Furthermore, the likelihood of rapid implementation and scalability is increased if the diverse actors that will have to contribute to the implementations of a regime have also participated in its elaboration.
- **Issue-based**, because each topic involves different sets of concerned stakeholders, or even different individuals and units within each entity. Efficient policy innovation therefore requires focus on a specific issue to ensure inclusion of all relevant actors.

Figure 2: Filling the Institutional Gap



Source: Authors.

Figure 3: Six Stakeholder Groups



Source: Authors.

Based on the lessons of the Internet & Jurisdiction Project, some key factors for the success of such issue-based policy networks are:

- framing the problem as an issue of common concern for all actors;
- ensuring the neutrality of the convener and facilitation team/secretariat;
- involving all six stakeholder groups: states, Internet platforms, technical operators, academia, civil society, and international organizations (see Figure 3);
- engaging a critical mass of actors with sufficient diversity to be representative of the various perspectives and to implement potential solutions;

- constructing and expanding a global network of key actors;
- creating trust among heterogeneous actors and adopting a shared vernacular;
- combining smaller working groups and reporting on progress to make the process manageable and transparent;
- informing stakeholders about relevant trends around the world to foster evidence-based policy innovation; and
- providing sufficient geographic diversity from the onset to allow the scalability of adoption of any emerging policy solution.

Addressing jurisdiction issues on the Internet and preempting the current legal arms race requires enhanced efforts to catalyze multi-stakeholder cooperation on the specific topics of cross-border requests for domain seizures, content takedowns and access to user data.

TOWARD TRANSNATIONAL FRAMEWORKS

Such innovative multi-stakeholder networks can produce scalable and adaptive policy standards that guarantee procedural interoperability and transnational due process in relations between public and private actors.

PROCEDURAL INTEROPERABILITY

International human rights frameworks already represent an overarching substantive reference at the global level. Recent UN Human Rights Council (2014) resolutions have reaffirmed that they apply online as well as off-line. However, rapid substantive legal harmonization at a more detailed level regarding use of the Internet is unrealistic, given the diversity of legislations that are often considered strong elements of national identity. Meanwhile, cross-border requests for domain seizures, content takedowns and access to user data pose everyday problems that require urgent action, as the stakes involved are high.

In contrast to traditional interstate cooperation, these increasingly cross-border interactions engage heterogeneous public and private actors. They are conducted in all shapes and formats, through broadly diverse communication channels, and often without clear and standardized procedures or sufficient transparency. In that context, prioritizing the development of shared procedural standards has several benefits:

- It provides a field of cooperation that helps build trust among stakeholders and paves the way for constructive discussions on contentious substantive norms.

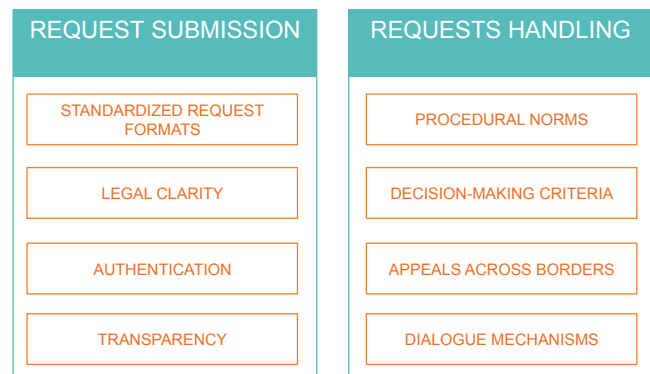
- It establishes interoperability among heterogeneous actors by providing shared vernacular and mechanisms for their interactions, not unlike the Transmission Control Protocol/IP enabled interoperability between heterogeneous networks.
- It prepares a future digitization of the request treatment workflow, in order to reduce the delays that plague current mechanisms, such as MLATs.
- Most important, it is an opportunity to incorporate due process requirements in operational frameworks by design, in order to improve transnational interactions and safeguard users' rights across borders.

TRANSNATIONAL DUE PROCESS

After four years of international multi-stakeholder discussions facilitated by the Internet & Jurisdiction Project, key elements of transnational due process have been identified with the goal of providing avenues for best practices, improving existing mechanisms such as MLATs and identifying a potential architecture for novel cooperation frameworks.

This architecture for transborder requests deals with two aspects: how requests are submitted and how requests are handled (Figure 4).

Figure 4: Architecture for Transnational Due Process Frameworks



Source: Authors.

The submission of requests raises the following sets of questions:

- How can request formats be standardized? What are current best practices? How can we incorporate due process by design into such formats?
- How can we ensure legal clarity for both intermediaries — potentially subjected to 190-plus different jurisdictions — and for users who struggle to understand the rights and obligations that apply to them in cyberspace?

- How can we build trust between senders and recipients of cross-border requests through authentication, in order to avoid abuses and arbitrary requests?
- What are best practices for transparency reporting? How can we spread this practice among public and private actors to increase accountability?

How requests are handled addresses the following components:

- What procedural norms must be respected by senders and recipients to make requests legitimate?
- Which decision-making criteria can ensure the respect of human rights and guarantee proportionality?
- What procedures can allow affordable and efficient redress by parties, especially users, across borders?
- How can trusted and efficient communication channels be constructed across borders to mitigate escalating tensions between public and private actors, especially in cases of non-compliance with requests?

While each of these questions can be further broken down into sub-elements, they will not be described here, as the above list is intended principally as a framework for discussions.

GOVERNANCE THROUGH POLICY STANDARDS

Norms and procedures developed through such multi-stakeholder processes can be considered “policy standards.” As innovative transnational cooperation frameworks, they can establish mutual commitments between the different stakeholders, with:

- clear distribution of responsibilities;
- specific norms, procedural mechanisms or guarantees; and
- clear decision-making criteria.

As new forms of transnational soft law, such operational governance frameworks can, in the context of addressing jurisdiction on the Internet, guarantee procedural interoperability and due process. In doing so, they can either help to reform existing modes of interstate cooperation (for example, the MLAT system) or fill current governance voids that require new sets of norms and standards.

Implementation and enforcement of such policy standards can leverage a combination of existing tools and cover the range from simple best practices to strict normative obligations. Public and private actors have different options to operationalize these shared norms through measures such as states referencing policy standards in

their administrative procedures, or Internet platforms and technical operators doing so in their terms of service. Multi-stakeholder policy standards can even be institutionally embedded in national laws, endorsed by international organizations or enshrined in new international treaties.

Drawing lessons from the governance “of” the Internet, a major advantage of standards is their potential to scale. Multi-stakeholder policy standards are based on consensus among different stakeholder groups, which augments the likelihood of successful and efficient adoption. They can more easily be implemented across heterogeneous public and private governance systems, which is the key to creating interoperability. Moreover, such policy standards can be improved and adapted more quickly than conventional treaties, which allows them to develop further as the Internet ecosystem evolves.

CONCLUSION

Thomas Kuhn, in his *Structure of Scientific Revolutions* (1962), describes paradigm shifts that modify the model underpinning a particular field when it no longer reflects or correctly explains observations. The Copernican revolution in astronomy is the most familiar example, triggered by the observations of Galileo’s telescope. Similarly, political paradigm shifts occur when a particular model of societal organization struggles to adequately address all problems of the time.

Rooted in the treaties of the Peace of Westphalia of the seventeenth century, our international system, based on the territorial jurisdictions, the separation of sovereignties, and non-interference, struggles to handle the transborder digital realities of the twenty-first century. The Internet acts like Galileo’s telescope, showing that traditional principles and approaches can become as much an obstacle as a solution to address the jurisdiction challenge in cross-border online spaces.

Addressing issues related to governance “on” the Internet requires a paradigm shift: from international cooperation only between states, to transnational cooperation among all stakeholders; from pure intergovernmental treaties to policy standards; and from intergovernmental institutions to issue-based governance networks.

Far from a rejection of traditional international cooperation, however, this is proposed as a constructive extension — a way to look at current practices in a new, generalized light. In physics, two theories coexist at the same time: relativity theory applies at high velocities in space; but in normal conditions, classical Newtonian, equations still allow us to build bridges and predict trajectories. Both have their respective zones of validity. Likewise, the type of transnational cooperation envisioned here in no way suppresses or reduces the relevance and authority of existing governance frameworks, in particular

national governments. On the contrary, multi-stakeholder processes can produce policy standards that inform the reform of existing interstate cooperation mechanisms, and policy standards can even later be enshrined by traditional multilateral organizations.

The global community needs to step up efforts to avoid the negative consequences of a legal arms race, preserve the global nature of the Internet and address its misuse. We need innovative cooperation mechanisms that are as transnational as the Internet itself and the necessary policy networks and ongoing dialogue processes to produce them.

WORKS CITED

- Boston Consulting Group. 2014. *Greasing the Wheels of the Internet Economy*. www.icann.org/en/system/files/files/bcg-internet-economy-27jan14-en.pdf.
- Council of Europe. 2011. *Recommendation CM/Rec (2011)8 of the Committee of Ministers to Member States on the Protection and Promotion of the Universality, Integrity and Openness of the Internet*.
- . 2015. *Recommendation CM/Rec(2015)6 of the Committee of Ministers to Member States on the Free, Transboundary Flow of Information on the Internet*. <https://wcd.coe.int/ViewDoc.jsp?id=2306649>.
- Cutler, A. Claire. 2001. "Critical reflections on the Westphalian assumptions of international law and organization: a crisis of legitimacy." *Review of International Studies* 27 (02): 133–50.
- de La Chapelle, B. 2007. "The Internet Governance Forum: How a United Nations Summit Produced a New Governance Paradigm for the Internet Age." In *Governing the Internet: Freedom and Regulation in the OSCE Region*, edited by OSCE, 27.
- De Werra, Jacques. 2015. "Alternative Dispute Resolution in Cyberspace: Can ADR Address the Challenges of Massive Online Micro Justice?" Presentation at the University of Geneva, November 27. http://svir-ssdi.ch/fileadmin/user_upload/VR-Tage/SSDI_-_Jde_Werra_-_ADR_24_11_2015_.pdf.
- Drake, W. J., V. G. Cerf and W. Kleinwächter. 2016. *Internet Fragmentation: An Overview*. World Economic Forum Future of the Internet Initiative White Paper. Geneva, Switzerland: World Economic Forum. www3.weforum.org/docs/WEF_FII_Internet_Fragmentation_An_Overview_2016.pdf.
- Facebook. .2015. Facebook Transparency Report. Government requests report January — June 2015. <https://govtrequests.facebook.com/>.
- Fehlinger, P. 2014. "Cyberspace Fragmentation: An Internet Governance Debate beyond Infrastructure." *Internet Policy Review*. <https://policyreview.info/articles/news/cyberspace-fragmentation-internet-governance-debate-beyond-infrastructure/266>.
- Google. 2015. Google Transparency Report. Requests for user information January — June 2015. www.google.com/transparencyreport/userdatarequests/countries/.
- ICANN. 2013. "Montevideo Statement on the Future of Internet Cooperation." www.icann.org/news/announcement-2013-10-07-en.

- . 2014. *Towards a Collaborative, Decentralized Internet Governance Ecosystem: Report by the Panel on Global Internet Cooperation and Governance Mechanisms*. Retrieved from www.internetsociety.org/sites/default/files/Internet%20Governance%20Report%20iPDF.pdf.
- Internet & Jurisdiction. n.d. Internet & Jurisdiction Retrospect Archive. www.internetjurisdiction.net/observatory/.
- Internet & Jurisdiction Retrospect. 2012a. "US authorities seize foreign .com gambling site registered in Canada via VeriSign." www.internetjurisdiction.net/observatory/retrospect/2012-february/.
- . 2012b. "US authorities give back seized domains of Spanish Rojadirecta site." www.internetjurisdiction.net/observatory/retrospect/2012-august/.
- . 2012c. "India cracks down on online content that stirs violence in its jurisdiction, needs US assistance to trace origins." www.internetjurisdiction.net/observatory/retrospect/2012-august/.
- . 2012d. "Pakistani YouTube block remains intact due to absence of MLAT with US jurisdiction." www.internetjurisdiction.net/observatory/retrospect/2012-november/.
- . 2013. "US Sponsored Stories Facebook settlement triggers Terms of Service changes for global users." www.internetjurisdiction.net/observatory/retrospect/2013-august/.
- . 2014a. "Hong Kong DPA wants to extend European de-index right on Google to Asia-Pacific region." www.internetjurisdiction.net/observatory/retrospect/2014-june/.
- . 2014b. "Marco Civil puts Brazilian data stored abroad under Brazilian jurisdiction." www.internetjurisdiction.net/observatory/retrospect/2014-april/.
- . 2014c. "Twitter blocked in Turkish jurisdiction at IP level." www.internetjurisdiction.net/observatory/retrospect/2014-march/.
- . 2015a. "Microsoft appeals US court order to hand over data stored in Ireland to US law enforcement." www.internetjurisdiction.net/observatory/retrospect/2015-september/.
- . 2015b. "New privacy standards: EU agrees on final draft of its data protection reform." www.internetjurisdiction.net/observatory/retrospect/2015-december/.
- . 2015c. "Belgium asserts jurisdiction over Yahoo, refuses MLAT procedure." www.internetjurisdiction.net/observatory/retrospect/2015-december/.
- . 2015d. "Brazil blocks WhatsApp for 12 hours with accidental impacts in Venezuela and Chile." www.internetjurisdiction.net/observatory/retrospect/2015-december/.
- . 2015e. "French DPA rejects Google's appeal on global application of 'right to be de-indexed.'" www.internetjurisdiction.net/observatory/retrospect/2015-september/.
- . 2015f. "EUDPAs meet with Google, Microsoft, Yahoo to discuss 'right to be de-indexed.'" www.internetjurisdiction.net/observatory/retrospect/2014-july/.
- . 2015g. "Constitutional Court of Colombia rules on 'right to be de-indexed' case." www.internetjurisdiction.net/observatory/retrospect/2015-july/.
- . 2015h. "Russian data-localization law might be delayed, some firms could be exempted." www.internetjurisdiction.net/observatory/retrospect/2015-july/.
- . 2016. "UK Home Office reaffirms the extraterritorial reach of the Draft Investigatory Powers Bill." www.internetjurisdiction.net/observatory/retrospect/2016-january/.
- Internet Society. n.d. The Internet and the Public Switched Telephone Network. www.internetsociety.org/sites/default/files/The%20Internet%20and%20the%20Public%20Switched%20Telephone%20Network.pdf.
- Kleinwächter, Wolfgang. 2001. "Global Governance in the Information Age." Papers from the Centre for Internet Research. http://cfi.au.dk/fileadmin/www.cfi.au.dk/publikationer/cfis_skriftserie/003_kleinwachter.pdf.
- Knill, C and D. Lehmkuhl. 2002. "Private Actors and the State: Internationalization and Changing Patterns of Governance." *Governance* 15 (1): 41.
- Kuhn, T. 1962. *The Structure of Scientific Revolutions*. Chicago, IL: University of Chicago Press.
- McCullagh, Declan. 2010. "State Dept. presses China ambassador on Google." CNET, January 22. www.cnet.com/news/state-dept-presses-china-ambassador-on-google/.
- Microsoft. 2015. Microsoft Transparency Hub. Law Enforcement Requests Report January — June 2015. www.microsoft.com/about/corporatecitizenship/en-us/transparencyhub/lerr/.

- NETmundial. 2014. "Roadmap Section IV: Jurisdiction Issues and How They Relate to Internet Governance." In *NETmundial Multistakeholder Statement*. <http://netmundial.br/wp-content/uploads/2014/04/NETmundial-Multistakeholder-Document.pdf>.
- NETmundial Initiative. n.d. "Terms of Reference." www.netmundial.org/terms-reference.
- Nye, Joseph S., Jr. 2014. *The Regime Complex for Managing Global Cyber Activities*. Global Commission on Internet Governance Paper Series No. 1. Waterloo, ON: CIGI.
- Ostrom, E. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge, UK: Cambridge University Press.
- Palfrey, John and Urs Gasser. 2012. *Interop: The Promise and Perils of Highly Interconnected Systems*. New York, NY: Basic Books.
- Raymond, M. and L. DeNardis. 2015. "Multistakeholderism: Anatomy of an Inchoate Global Institution." *International Theory* 7 (3): 572–616.
- Scassa, Teresa and Robert J. Currie. 2010. "New First Principles: Assessing the Internet's Challenges to Jurisdiction." *Georgetown Journal of International Law* 42 (4): 1018.
- Twitter. 2015. Twitter Transparency Report. Information requests January — June 2015. <https://transparency.twitter.com/information-requests/22015/jan-jun>.
- UNESCO. 2013. "Internet Universality: A Means Towards Building Knowledge Societies and the Post-2015 Sustainable Development Agenda." UNESCO Discussion Paper. www.unesco.org/new/fileadmin/MULTIMEDIA/HQ/CI/CI/pdf/news/internet_universality_en.pdf.
- UN Human Rights Office of the High Commissioner. 2011. *Report of the Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression, Frank La Rue*. www2.ohchr.org/english/bodies/hrcouncil/docs/17session/a.hrc.17.27_en.pdf.
- . 2015. *Report of the Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression, David Kaye*. www.ohchr.org/EN/HRBodies/HRC/RegularSessions/Session29/Documents/A.HRC.29.32_AEV.doc.
- UN Human Rights Council. 2014. *Resolution A/HRC/26/L.24 on the Promotion, Protection and Enjoyment of Human Rights on the Internet*. <http://daccess-dds-ny.un.org/doc/UNDOC/LTD/G14/059/67/PDF/G1405967.pdf?OpenElement>.
- Weber, Rolf H. 2014. *Legal Interoperability as a Tool for Combatting Fragmentation*. Global Commission on Internet Governance Paper Series No. 4. Waterloo, ON: CIGI.
- WSIS. 2005. "Tunis Agenda for the Information Society." WSIS-05/TUNIS/DOC/6(Rev. 1)-E, November 18. www.itu.int/net/wsisis/docs2/tunis/off/6rev1.html.
- Yahoo. 2015. Yahoo Transparency Report. Government data requests January — June 2015. <https://transparency.yahoo.com/government-data-requests/index.htm>.

ABOUT THE AUTHORS

Bertrand de La Chapelle is co-founder and director of the Internet & Jurisdiction Project. He was a director on the Internet Corporation for Assigned Names and Numbers (ICANN) Board from 2010 to 2013. From 2006 to 2010, he was France's Thematic Ambassador and Special Envoy for the Information Society. In this position, he participated in all World Summit on the Information Society (WSIS) follow-up activities and Internet governance processes, including in particular the Internet Governance Forum, and was a vice-chair of ICANN's Governmental Advisory Committee. Between 2002 and 2005, he actively participated in the WSIS to promote dialogue among civil society, the private sector and governments. He is a graduate of Ecole Polytechnique, Sciences Po Paris and Ecole Nationale d'Administration.

Paul Fehlinger is co-founder and manager of the Internet & Jurisdiction Project. He is actively engaged in global Internet governance fora, speaking at venues such as the Internet Governance Forum, the Organisation for Economic Co-operation and Development, and The Council of Europe. Paul was appointed to the Advisory Network of the Global Commission on Internet Governance and to the Working Group on Rule of Law of the Freedom Online Coalition. He is also a participant in the Council of Europe Committee of Experts on Cross-border Flow of Internet Traffic and Internet Freedom, and the World Economic Forum's Future of the Internet Initiative. He holds a master's degree in international relations from Sciences Po Paris, was a scholar of the German National Merit Foundation and has also worked in journalism.

CHAPTER SIX: LEGAL INTEROPERABILITY AS A TOOL FOR COMBATING FRAGMENTATION

Rolf H. Weber

Copyright © 2014 by Rolf H. Weber

INTRODUCTION

INTEROPERABILITY IN GENERAL

The term interoperability is commonly understood in the infrastructure context, namely as a tool to interconnect networks. In general, open standards and interoperable systems make life easier and increase efficiency. Interoperability functions can be identified on four broad layers of complex systems (Palfrey and Gasser 2012, 5-6): technology — the ability to transfer and render data and other information across systems, applications or components; data — the ability to read the data; human elements — the ability to communicate, for example, through a common language; and institutional aspects — the ability to work together (Weber 2014, 143).

In a broad sense, conditions for non-restricted interoperability can encompass access to the decision-making processes, transparent and undistorted procedures, pro-competitive goals, objective and relevant criteria for technology selection, and renunciation of over-standardization (Brown and Marsden 2013, 28-29). In a narrow sense, interoperability between networks refers to the possibility of easily linking different legal structures; insofar, a too low level of interoperability leads to a non-optimal level of interconnectedness.

An open and interoperable environment can stimulate innovation since state censorship and private control of general value chains might make innovation difficult; the wider the choice available to users, the higher their ability to take advantage of their freedoms, even without a guarantee of fundamental rights (Weber 2014, 144). Usually, a combination of legal instruments is needed to reach optimal levels of interoperability in practice, depending on the applied or developed architecture; for example, cloud computing, smart grids or the Internet of Things (Palfrey and Gasser 2012, 160, 232–51).

From a theoretical perspective, interoperability issues can be mapped by differentiating between government-driven measures and private-sector-led approaches on the one hand, and unilateral and collaborative approaches on the other (*ibid.*, 14). Governmental actions encompass the disclosure of information, a transparency regime or public procurement rules; private initiatives include reverse engineering, licensing, technical collaboration and open standards initiatives (Weber 2014, 144). In a layer model, legal operability must be put into the appropriate relation to other layers, for example, the organizational, semantic and technical layers (European Commission 2010, 21).

LEGAL INTEROPERABILITY

Legal interoperability addresses the process of making legal rules cooperate across jurisdictions, on different subsidiary levels within a single state or between two or

more states. Whether new laws should be implemented or existing laws adjusted or reinterpreted to achieve this interoperability depends on the given circumstances (Palfrey and Gasser 2012, 178-79). In view of the increasing fragmentation of the legal environment in cyberspace, efforts must be undertaken to achieve higher levels of legal and policy interoperability in order to facilitate global communication, to reduce costs in cross-border business, and to drive innovation and economic growth. Interoperable legal rules can also create a level playing field for the next generation of technologies and cultural exchange (Weber 2014, 153; Gasser and Palfrey 2012, 132-33).

This chapter examines the rising debate of legal interoperability and discusses the different regulatory models available in order to make legal rules interoperable. Theoretically, legal interoperability can be looked at from the angles of substance or procedure. This chapter focuses on the issue of substantive or normative concerns and does not address procedural structures in detail (for example, legal jurisdiction or multi-stakeholder participation).¹

The degree of legal interoperability depends on the material issue at stake. For example, harmonized legal rules are important for the implementation of the Domain Name System (DNS); however, less unification appears to be needed in the field of cultural expression. Therefore, this chapter addresses the following questions:

- What relevance and facets does legal interoperability have in the context of Internet governance?
- How should a matrix of the available regulatory models be designed in the Internet governance framework, and which segments of regulatory intervention could be distinguished?
- How can substantive legal interoperability be used as a tool to combat fragmentation?

CHARACTERISTICS AND IMPORTANCE OF LEGAL INTEROPERABILITY

LEGAL INTEROPERABILITY AS A NORMATIVE TOOL

The supranational realization of the process of legal interoperability (as the process of making legal rules work together across jurisdictions) can fluctuate between two poles: full harmonization and a complete fragmentation on a bilateral, plurilateral or multilateral level. In the first scenario, all laws would be the same everywhere; in the second, the legal systems would be so different in each country that economic, social and cultural

¹ For further details on multi-stakeholder participation, see Weber (2014, 126–35).

interactions become impossible (Palfrey and Gasser 2012, 181). Obviously, the two extremes do not correspond to reality, as the law does not reflect them; depending on the circumstances, the ideal is usually between the two poles, i.e., closer to harmonization or to fragmentation as required by practical considerations. In addition, public policy issues can play a role (European Commission 2010, 22). An in-between level of legal interoperability can usually be considered as good policy (Palfrey and Gasser 2012, 184).

In order to give some guidance to the applicable normative system, the legal community has developed rules on conflicts of law. These rules help determine which legal system should be applied in a given case. However, the rules on conflicts of law (private international law) do not overcome the substantive differences in national legal orders (and therefore do not lead to legal interoperability), they *only* give guidance on how to determine the applicable normative rules. This assessment does not mean that such rules do not have any impact on the substantive contents of interoperable legal systems, but their influence is of an indirect nature. As a consequence, venue selection by the parties and public interest exceptions to such selection gain practical importance.

The normative objective of legal interoperability consists of the attempt to combat legal fragmentation caused by different national law systems. However, national rules are a consequence of the sovereignty principle and, therefore, are legitimate to the extent of the justified scope of sovereignty (Weber 2014, 7–12). In addition, the more legal interoperability is achieved, the narrower the scope of legal competition between nation-states will be; consequently, a fragile equilibrium must be balanced out.

ADEQUATELY STRUCTURED DEGREES OF LEGAL INTEROPERABILITY

The relationship between law and interoperability must be understood as a multidirectional network. Legal interoperability should make systems work together, but not make the systems all the same, since regulatory competition can be advantageous and productive provided the best normative order prevails (Palfrey and Gasser 2012, 179). Furthermore, changes in the legal order and/or in the interoperability regime have an impact on the design of the relationship between the two.

Higher levels of legal interoperability usually require a more careful design of governmental regulations and a disclosure of the rules in order to increase legal certainty (ibid., 178). The highest level would be reached in the case of a total harmonization of normative rules. However, a total harmonization should not be the approach to follow in all cases since, on the one hand, such a framework could not take into account the cultural diversity of societies in the global online world and, on the other, would be a

utopian wish in reality. Moreover, it is important to find the appropriate degree of legal interoperability (instead of an all-or-nothing solution) considering the substantive principles (such as freedom of expression or privacy) in different circumstances.²

Consequently, legal operability is a matter of degrees (ibid., 183): a very high level of legal interoperability could cause difficulties in the application of the harmonized rules on the national level (for example, due to the difficulty of reaching harmonized interpretation methods), while a very low level of legal interoperability could provoke challenges in respect of the smooth (social or economic) interaction (ibid., 2). As in the case of the appropriate level of interconnectedness, the rule makers have to find the optimal degree of legal interoperability.

COST-BENEFIT ANALYSIS OF LEGAL INTEROPERABILITY

In the information society in particular, legal interoperability drives innovation, competition, trade and economic growth (Palfrey and Gasser 2012, 182); furthermore, costs associated with doing business across borders are reduced. This assessment can be seen in the example of non-founding countries entering the World Trade Organization (WTO) in 1994. These countries are usually obliged to large-scale changes in many business laws relevant for international trade as negotiated in the so-called “accession protocol.”³ Even if total harmonization is not envisaged and regularly not achieved, an increased degree of legal interoperability as acknowledged by a WTO applicant facilitates cross-border trade.

Besides economic factors, higher levels of legal interoperability can also help secure freedom of expression and foster diversity of other fundamental rights, as well as lead to better laws (ibid., 179, 183). This function is mainly realized by international organizations such as the United Nations or the Council of Europe.⁴ An example of this is the prohibition of child labour as stated by the UN Convention on the Rights of the Child.⁵ As far as the freedom of expression is concerned, Internet service

2 For further details, see the Case Studies section on page 10.

3 The protocols for new members since 1995, including commitments in goods and services, are available at www.wto.org/english/thewto_e/acc_e/completeacc_e.htm.

4 See, for example, the Council of Europe’s Declaration by the Committee of Ministers on Internet Governance Principles, which invites its member states to comply with basic online fundamental freedoms by, among others, referring to the protection of all fundamental rights (principle 1), the responsibilities of states (principle 3) or the empowerment of Internet users (principle 4), available at <https://wcd.coe.int/ViewDoc.jsp?id=1835773>.

5 Additional obligations in connection with child labour are contained in various International Labor Organization declarations.

providers (ISPs) have gained increased legal certainty on a regional level by way of the E-Commerce Directive implemented by the European Union in 2000.⁶

IMPLEMENTATION OF LEGAL INTEROPERABILITY

Legal interoperability can be implemented by applying a top-down model or a bottom-up process. As far as the intensity of achieving legal interoperability is concerned, a distinction between harmonization, standardization, mutual recognition and other approaches is possible.

TOP-DOWN AND BOTTOM-UP APPROACHES

A top-down approach necessarily requires the establishment of a global agency, for example, the United Nations or any of the UN special organizations. Usually, such an approach generates the implementation of large bureaucracies (Palfrey and Gasser 2012, 184-85). In the context of Internet governance, the International Telecommunication Union (ITU) appears to be the most prominent top-down actor; however, as the experience of the World Conference on International Telecommunications in Dubai showed in December 2012, the attempt to agree by consensus on new rules, not even directly related to Internet governance, failed and common visions of global norm-setting did not evolve (Weber 2014, 102-03).

A bottom-up process to achieve legal interoperability must be based on a step-by-step model that encompasses the major concerned entities and persons of the substantive topic (Palfrey and Gasser 2012, 185). The NETmundial, held in Sao Paulo in April 2014, embodied a relatively successful bottom-up process, wherein the various stakeholders are principally granted equal rights in the negotiation processes of the final non-binding declaration.⁷ The Global Network Initiative, which encompasses major Internet and information technology companies, can also be seen as a bottom-up model. Generally speaking, a bottom-up approach requires a large amount of coordination, but no harmonization or management by central bodies; thereby, coordination processes can be time-consuming and somewhat cumbersome.

6 This is the “Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular electronic commerce, in the Internal Market (‘Directive on electronic commerce’).” See <http://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32000L0031>.

7 “NETmundial Multistakeholder Statement,” April 24 2014, available at <http://netmundial.br/wp-content/uploads/2014/04/NETmundial-Multistakeholder-Document.pdf>.

REGULATORY MODELS AIMING AT LEGAL INTEROPERABILITY

Harmonization

Regulatory harmonization (a pillar of legal interoperability) can generally be defined as the legal model for institutionalizing a desired cooperation by confining actors and policies into the normative corset of rights and obligations (Weber 2009, 651, 658). Harmonization depicts the process of the unification of law, which often follows a previous approach of standardization. Therefore, harmonization should not be qualified as a contrast to standardization, but rather as a further step in the direction of legal convergence (ibid., 659). Harmonization can emerge in different degrees; for example, EU directives do not prescribe specific wordings for national legislation, but certain results that need to be achieved.

Harmonization as an objective does not necessarily define the type of national law that is employed. Moreover, on the basis of a cost-benefit analysis of the different forms of regulations, the choice must be made which regulatory technique is best suited for which type of legal issue (ibid.). The regulatory concept of harmonization also involves critical issues, one of which is the unification of the many existing national regulatory models. In practice, the choice is often made for the benefit of the legislation and regulatory practices of the most dominant state, which might contrast a large part of the global community. If the whole global community is involved in the preparation of harmonizing laws, there is a significant risk for a regulatory race to the bottom, as long as there is no need to tackle a duly acknowledged factual problem. If regulatory harmonization takes place on a relatively low level and in a generalized manner (an effect of the “highest common denominator”), the rules leave space for creative individual interpretation and compliance, which, in turn, leads to legal uncertainty (Weber 2009, 659).

Standardization

Standardization is usually defined as a regulatory approach that is based on widely accepted good principles, practices or guidelines in a given area; standards may also relate to the usual behaviour of the “reasonable man” (Miller 2007). Three types of standardization can be distinguished: technical, economic and legal. Technical standardization leads to technical interoperability. Economic standardization means that sellers would offer more interchangeability of their products than what is necessary and legally required. Legal standardization can be defined as an understanding approved by a recognized body that provides for common and repeated application, usually in the form of rules or guidelines. Mostly, legal standards express or stand for a general direction or a behavioural value, with which the average human or commercial entity is expected to comply. In order for a

standard to be effective, it is necessary that it addresses the concerned persons on all levels of business activities (Weber 2009, 660).

Standardization constitutes an important element in the process of regulating certain ways of behaviour: on the one hand, standardization encompasses the notion of making coherent, diverging technical characteristics; on the other hand, many standards qualify as soft law (Weber 2014, 22–32) that, even if lacking a legitimate authority for adoption and enforcement, provide a concrete and normatively relevant benchmark for the behaviour of the concerned community. Insofar, standardization can be seen as a first step to a later harmonization.

An important role in the context of standardization is played by standard-setting organizations (SSOs) developing international standards. Most SSOs are established as private entities (for example, as associations) and composed of national standards bodies; in the cyberspace field, the ITU is an exception as a treaty-based organization established as a permanent agency of the United Nations and driven by national governments as the primary members. In the Internet world, the most prominent SSOs are the Internet Engineering Task Force and the World Wide Web Consortium. The development of technical standards is usually concerned with interface standards making different systems interoperable; nevertheless, it cannot be overlooked that in many areas of technology, rigorous competition exists between different SSOs vying for leadership. SSOs can also contribute to the legal interoperability of contractual provisions and terms of service (Tamm Hallström and Boström 2010).

Mutual Recognition

Mutual recognition originally involved the assessment of comparability or equivalence of regulatory measures. Later, this assessment was converted into an independent legal principle.⁸ Put more simply, mutual recognition is the consent to compromise a country's regulatory autonomy by it accepting that another state's regulation is "good enough" or satisfactory; in other words, mutual recognition acknowledges that different national requirements can be interchangeable in order to be domestically applied (Weber 2009, 661–62). The principle of mutual recognition is widely accepted as a cross-border rule based on the concept that, even in the absence of harmonization, the foreign state has applied its norms with diligence and precaution, making them adequate for domestic application elsewhere.

Mutual recognition plays a crucial role in the European Union, where the "single passport" system within the region requires the need for privileged or facilitated market

⁸ Based on the decision of the European Court of Justice in *Cassis de Dijon*, *Rewe-Zentral AG v. Bundesmonopolverwaltung für Branntwein*, Judgment of the Court of 20 February 1979, Case 120/78.

access across borders (*ibid.*, 662). On a global level, the WTO's General Agreement of Trade in Services partly relies on the principle of mutual recognition, for example, with financial services. However, mutual recognition should be considered a second-best solution after harmonization or standardization, if legal interoperability is not achieved.

Other Approaches for Legal Interoperability

Reciprocity: This is a traditional principle in international law that attempts to achieve equilibrium between two countries regarding certain legal aspects. It generally refers to the balance of concessions to be sought in cross-border negotiations. Reciprocity is due to the commitments undertaken bilaterally if and to the extent agreed by the concerned parties. More recently, however, states are reluctant to apply reciprocity since this model only encompasses a narrow scope of legal interoperability and might also violate the most-favoured-nation principle in international instruments, for example in the context of the WTO.

Cooperation: In order to overcome the disparity of different legal regimes, regulators partly settle their responsibility by defining clear mandates and by agreeing on cooperation among themselves. Cooperation between different agencies can manifest in collective regulatory rules or at least lead to the agencies coordinating their efforts in designing, applying and enforcing different regulatory issues (Weber 2009, 664). But this approach is rather individualistic and often spontaneous. Based on the circumstances, agencies try to find an adequate solution to the occurring problem. This approach can make sense in a particular situation; however, cooperation does not contribute to an improvement of legal interoperability.

MAPPING OF REGULATORY MODELS

Mapping the different regulatory models and sources of law is outlined in Table 1, which contains some of the legal instruments available in the Internet governance context.

This table should be viewed in light of the substantive topics of Internet governance⁹ and needs further elaboration, even if the allocation of functions and activities is difficult to establish due to social and cultural perceptions. It can be stated that legal interoperability would be increased if substantive topics can be moved up and to the left. Nevertheless, it must be taken into account that lower-level arrangements that are actually applied and enforced can be more efficient than unexecuted higher-level theoretical models.

A method to potentially address the issue of "adequate" or "optimum" levels of legal interoperability could be to apply different regulatory models and mechanisms (according

⁹ For an overview of the topics see DeNardis and Raymond (n.d., 11–12).

Table 1: Normative Sources and Regulatory Concepts

Source of Law	Regulatory Models				
	Harmonization	Standardization	Mutual recognition	Reciprocity	Cooperation
Treaty law	ITU		EU E-Commerce Directive		Council of Europe Cybercrime Convention
Customs/standards		IETF technical standards			
General principles	Human rights declarations or recommendations			No-harm principle between states	
Self-regulation	ICANN DNS, Global Network Initiative	ISPs' codes of conduct		Data protection framework for business entities	ISPs' codes of conduct

Source: Author.

Note: Blank squares indicate that there is no instrument available.

to the given circumstances) that can enable, based on past experience, certain levels of legal interoperability within certain contexts. Consequently, the assessment of the degree and scope of legal interoperability, as well as its method of approach, depends on the substantive topic at hand. In order to illustrate this theoretical assessment, two case studies on freedom of expression and data protection principles are presented below to examine how the requirements of legal interoperability could be fulfilled.

PROCEDURAL ISSUES

As previously mentioned, legal interoperability is mainly an issue of cross-border coherence of normative orders, but procedural aspects can also play a role. The venue selection allows parties to choose the preferred normative order; venue selection is limited by public interest exceptions that restrict this choice and give a prevailing force to a specific national law. The venue selection can lead to legal interoperability within a private group, in the sense that all group entities are acting on the basis of the same normative order.

Another issue concerns the dispute resolution requirements. Depending on the resolution mechanism, a higher level of acceptance to a newly established substantive normative order can be achieved. The term “dispute resolution mechanism” should be understood broadly, including not only traditional proceedings, such as arbitration, but also all conceivable forms of mediation (Weber 2014, 148). Arbitration has reached legal interoperability due to the fact that enforcement of arbitral awards is possible according to the provisions of the 1958 New York Convention. New forms of alternative dispute resolution mechanisms should be taken into account, however, if the binding effects of norms can be achieved in the given circumstances. Dispute resolution mechanisms can be necessary to clarify which legal obligations are

potentially incomplete or inadequate. For example, even if a suitable forum for complaints in cyberspace is not yet available, consideration should be given to the implementation of new structures dealing with the settlement of the disputes (Weber 2012, 9-10).

CASE STUDIES

From a conceptual perspective, five major features of global Internet governance can be distinguished: the arrangements of the technical architecture, the Internet governance infrastructure, the privatization of governance mechanisms, the Internet control points as sites of global conflict and the regional initiatives addressing geopolitical strategies (DeNardis 2014, 7–19). In other words, Internet governance encompasses the design and administration of the technologies necessary to keep the Internet operational and the enactment of substantive policies around these technologies (ibid., 6). From this broad array of issues, many examples could be chosen for an elaboration of the strengths and weaknesses of legal interoperability,¹⁰ but the cases discussed here are freedom of expression and the data protection framework.

FREEDOM OF EXPRESSION

A challenging topic in the context of legal interoperability is the conciliation of the different understandings of, and the manifold cultural approaches to, freedom of expression. Freedom of expression is a fundamental right that is acknowledged in many international instruments (such

¹⁰ For an overview of issues, see the respective list published by the Berkman Center of Harvard Law School, available at <http://cyber.law.harvard.edu/research/interoperability>. A practical and important topic concerns the license interoperability; for further details see Morando (2013). Difficult questions also arise in connection with cyber security, these issues being a particularly sensitive area of achieving legal interoperability (see Palfrey and Gasser 2012, 188-89).

as the United Nations and through regional conventions), but the provisions often contain a reservation allowing the implementation of state legislation based on the principle of public order. Interpretation of this reservation is subject to social and cultural perceptions, and therefore legal interoperability is unlikely to be achieved. For example, the likelihood of the First Amendment to the United States Constitution (which includes the freedoms of religion, speech, the press and association) becoming the rule in China or the Middle East is extremely low (Palfrey and Gasser 2012, 181). However, even if the cultures of societies involved in cross-border activities are relatively similar (such as with Europe and the United States), substantial problems can occur. The most famous cases dealing with freedom of expression were *Ligue contre le racisme et l'antisémitisme et Union des étudiants juifs de France (LICRA) v. Yahoo!*, and *Google and the right to be forgotten*.

The Case of Yahoo!

Yahoo! operated an auction business from its California base offering thousands of items of Nazi memorabilia for sale. LICRA, a French anti-racism and anti-Semitism organization, started legal action against Yahoo!, alleging that the company was violating French law by providing access to these materials through its website. Essentially, the French courts not only acknowledged their jurisdiction (competence) in a case against a US company, but also applied French law prohibiting a US firm from operating auctions that sell “critical” goods to French citizens in violation of French law.¹¹ Consequently, the freedom of advertising for some goods as emanation of the freedom of expression was restricted.

The Case of Google

In May 2014, the Court of Justice of the European Union (CJEU) requested that Google Spain remove a link providing information about a seizure of assets of a Spanish citizen some 15 years ago.¹² The decision was based on the EU Data Protection Directive being interpreted beyond its wording as containing a “right to be forgotten.”¹³ Google’s argument that the removal of the link would contradict the fundamental right of freedom of

expression as guaranteed by different international legal instruments as well as by the US Constitution did not convince the CJEU. On the contrary, the CJEU regarded the individual’s interest in removing links with personal information that is inadequate, irrelevant or excessive as being more important than the public’s interest in getting access to that kind of information. According to the CJEU, the economic interests of the search engine do not justify the interference with a person’s right to data protection; that is, the freedom of expression can be legitimately limited in the interests of privacy.

In Search of Alternatives: Codes of Conduct

These cases show that far-reaching legal interoperability can hardly be achieved by harmonization of law through international instruments. However, in this context, ISPs have the option to agree on codes of conduct standardizing intervention practices; contrary to a mandatory provision, codes of conduct do not legally oblige their addressees, but take full effect as voluntary self-regulation. A practical realization of this approach can be seen in the efforts of the Global Network Initiative attempting to incentivize Internet and IT companies to comply with some commonly accepted standards (such as freedom of expression or privacy).

Having been constructed by engineers, the Internet and its content, services and applications are based on technology rather than on legal instruments; anyone should be able to design new Internet content using publicly and freely available protocols and software (Brown and Marsden 2013, 7-8). Being a public common good that is based on the good conduct of its users, Internet pioneers since the beginning of the World Wide Web realized that most Internet functions required trust.¹⁴ Too much control from national or international legislators would impair the free development of the Internet, which in turn would be contrary to the Internet’s basic principle of being a global medium with an infinite spectrum and low barriers to entry.

Accordingly, self-regulation and minimal state involvement appear to be more efficient regulatory instruments to “regulate” the Internet than international treaties (ibid., 2). In other words, legal interoperability might be improved on the basis of ISPs’ codes of conduct containing rules in respect of the freedom of expression; however, this “improvement” also carries the risk that private actors are empowered to technically design the scope of a fundamental right. In Table 1, the most ideal approach seems to consist in standardization based on a self-regulatory regime.

11 The Tribunal de grande instance in Paris confirmed the illegal nature of the sale of Nazi-era memorabilia under French law in 2000 (thereby approving the competence of the French courts in a complaint against the US firm Yahoo!; decision RG:00/0538 of May 22, 2000 and November 22, 2000). Later, Yahoo! began legal action in the United States, arguing that the sale’s prohibition would contradict the First Amendment of the US Constitution.

12 *Google Spain SL and Google Inc. v Agencia Española de Protección de Datos (AEPD) and Mario Costeja González*, Judgment of the Court of May 13, 2014, case C-131/12.

13 Directive 95/46/EC of the European Parliament and of the Council of October 24, 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data.

14 See the Internet Society’s Internet Code of Conduct, available at www.isoc.org/internet/conduct/.

DATA PROTECTION FRAMEWORK

Privacy can be examined from technical and legal perspectives. The increase in technical interoperability raises concerns that it may make systems less secure. However, security problems are not related to interoperability as such but rather, in what interoperability makes possible (Palfrey and Gasser 2012, 77). As in the case of legal interoperability, the optimal degree of technical interoperability varies depending on the circumstances; consequently, engineers need to implement designs of selective interoperability or limited interoperability (*ibid.*, 79–80). The main emphasis in the following sections is on the legal interoperability of data protection rules.

In view of the massive growth in the complexity and volume of transborder data flows, accompanied by a change in the nature of such transfers, theoretically, global privacy rules should be available. In practice, however, data protection laws are very different in the various regions of the world, and a harmonization of these rules is not expected in the near future (Weber 2013, 1–3). The lack of harmonized global rules governing transborder data flows causes several risks: business challenges, particularly in outsourcing transactions; technological challenges in view of the growing data warehouses and increased data mining; and security challenges, since large data collections are a threat to security (Gunasekara 2009, 147, 154–63).

From a theoretical perspective, the harmonization of data protection standards would certainly facilitate the transborder flow of information. Globally, however, such an objective is not likely to be achieved, even if some progress has been made on the harmonization of rules on a regional level, for example, among EU member states. Additionally, pressure to harmonize data protection standards comes from international trade law: different levels of protection can jeopardize the cross-border rendering of services, particularly IT and electronic commerce services (Weber 2013, 5).

Due to the complexity of technology, such as cloud computing, regulations become difficult to implement and their enforcement is cumbersome. Therefore, regulations should enable individuals and businesses to reach a high level of compliance at a reasonable cost; besides, regulators are called on to design norms that are more efficient. In this context, transparency could facilitate the decision-making processes for businesses considering how to handle transborder data flows. Transparency could be increased by making all relevant texts of national laws and regulations on data protection, particularly on transborder data flows, available in different languages on the Internet; by providing regular and timely updates of the respective legal rules; and by designating a contact point in the government to which questions about transborder data transfers can be addressed (*ibid.*, 6).

With respect to the increase of transborder data flows, the “traditional” geographical approach of looking at the risks caused by the country or location to which the data are to be transferred no longer seems to be appropriate (*ibid.*). Moreover, an organizational model should be implemented that examines the risks caused by the holder (controller) of the data that are being transferred. This model would substantially address the fact that the data holder is responsible for the proper treatment of data when shipped abroad. Consequently, the organizational model burdens the data holder with the task of ensuring that the processing of data in another country is executed in accordance with the relevant data protection standards. This concept is based on the accountability principle: the appropriate level of protection is to be fulfilled by binding corporate rules or to be contractually designed by the parties involved (*ibid.*).

Consequently, and for good reasons, the draft for a new General Data Protection Regulation of the European Union puts a great deal of emphasis on binding corporate rules (BCR).¹⁵ In principle, each member of an organization has to sign the BCR, which should contribute to the realization of a minimum level of protection. The BCR standards must be internationally binding within the concerned organization; incorporate the material data processing rules; provide for a network of measures ensuring compliance with the applicable rules, for an internal complaints-handling process and for an auditing program; ensure suitable training for employees; and be enforceable by the beneficiaries of the BCR (Weber 2013, 12; Moerel 2012).

By acknowledging the validity of the BCR, the requirement that some principles must play an important role in data protection is realized; corporate law solutions as a self-regulatory mechanism can be a valid substitute for legislative measures and can establish a higher level of privacy than contested or ineffective multilateral treaty arrangements (Weber 2013, 12; Gunasekara 2009, 174–75). Table 1 shows the most ideal approach to be the reciprocity model based on a self-regulatory framework.

OUTLOOK

Legal interoperability is a very complex issue, and the costs of non-interoperable laws in a highly networked world will increase. The multiplicity of regulatory actors bears the risk of incoherent rule making; this risk is even enforced if regulatory actors try to expand their activities beyond their original mandate (Weber 2009, 682). Rule makers should, therefore, be smart about the design of law in view of the global information exchange (Palfrey and Gasser 2012, 177, 191, 256). Indeed, the exploration and

¹⁵ To view the unofficial consolidated version of the General Data Protection Regulation from June 28, 2014, see www.delegedata.de/wp-content/uploads/2014/06/DS-GVO-konsolidiert.pdf.

development of the substantive and structural dimensions of the nascent concept of legal interoperability (as a “third way” between fragmentation and harmonization) merit increased attention.

The key objective is the attempt to achieve interoperable rules that create a level playing field for the next generation of technologies and social exchange. If an adequate level of legal interoperability is not achieved and a far-reaching fragmentation prevails, the likelihood also increases that dominant states are inclined to enlarge the geographical scope of their laws by having them applied in an extraterritorial manner. This kind of legal harmonization would be to the detriment of non-dominant societies.

Currently, the efforts in analyzing the different available regulatory models’ strengths and weaknesses are in a state of infancy. For the time being, the traditional legal reality still consists of fragmentation, based on the sovereignty principle. This model must be changed, at least to a certain degree. Technical standardization and common understandings — with respect to generally applicable principles such as the no-harm, shared responsibility, good faith or ethical behaviour principles — need to be developed. Not every area of Internet governance needs the same level of harmonization, coordination or standardization.

Further research and thinking is needed. In particular, the procedural dimension of legal interoperability should also be explored, in addition to efforts toward its normative elements. In this regard, innovative operational approaches¹⁶ to legal cooperation are especially important whenever online interactions involve multiple jurisdictions at the same time and a convergence of laws is difficult to achieve. Different regulatory models that can serve the purposes of the manifold substantive topics are available, and nuances in the design of rule-making processes will gain importance. The unintended consequences of not having legally interoperable regimes must be avoided.

¹⁶ On this issue, see the work conducted within the Internet & Jurisdiction Project facilitated by Bertrand de La Chapelle and Paul Fehlinger in Paris, available at www.internetjurisdiction.net.

WORKS CITED

- Brown, Ian and Christopher T. Marsden. 2013. *Regulating Code: Good Governance and better Regulation in the Information Age*. Cambridge, MA: MIT Press.
- DeNardis, Laura. 2014. *The Global War for Internet Governance*. New Haven, CT: Yale University Press.
- DeNardis, Laura and Mark Raymond. n.d. “Multistakeholderism: Anatomy of an Inchoate Global Institution.” *Academia.edu*. www.academia.edu/9027904/Multistakeholderism_Anatomy_of_an_Inchoate_Global_Institution.
- European Commission. 2010. *Annex 2 to the Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of Regions: Towards Interoperability for the European Public Services*. Bruxelles: European Commission. http://ec.europa.eu/isa/documents/isa_annex_ii_eif_en.pdf.
- Gasser, Urs and John Palfrey. 2012. “Fostering Innovation and Trade in the Global Information Society: The Different Facets and Roles of Interoperability.” In *Trade Governance in the Digital Age*, edited by Mira Burri and Thomas Cottier, 123–53. Cambridge: Cambridge University Press.
- Gunasekara, Gehan. 2009. “The ‘Final’ Privacy Frontier? Regulating Trans-border Data Flows.” *International Journal of Law and Information Technology* 17 (2): 147–79.
- Miller, Alan D. 2007. “The ‘Reasonable Man’ and Other Legal Standards.” California Institute of Technology Social Science Working Paper No. 1277.
- Moerel, Lokke. 2012. *Binding Corporate Rules — Corporate Self-Regulation of Global Data Transfers*. Oxford: Oxford University Press.
- Morando, Federico. 2013. “Legal Interoperability: Making Open (Government) Data Compatible with Businesses and Communities.” *JLIS.IT* 4 (1): 441–52.
- Palfrey, John and Urs Gasser. 2012. *Interop: The Promise and Perils of Highly Interconnected Systems*. New York, NY: Basic Books (Persens Books Group).
- Tamm Hallström, Kristina and Magnus Boström. 2010. *Transnational Multi-Stakeholder Standardization*. Cheltenham, UK and Northampton, MA: Edward Elgar.
- Weber, Rolf H. 2009. “Mapping and Structuring International Financial Regulation — A Theoretical Approach.” *European Banking Law Review* 20 (5): 651–88.

- . 2012. “Future Design of Cyberspace Law — ‘Laws are Sand’ (Mark Twain, The Gorky Incident).” *Journal of Politics and Law* 5 (4): 1–14.
- . 2013. “Transborder Data Transfers: Concepts, Regulatory Approaches and New Legislative Initiatives.” *International Data Privacy Law*: 1–14. doi:10.1093/idpl/ipt001.
- . 2014. *Realizing a New Global Cyberspace Framework — Normative Foundations and Guiding Principles*. Zurich: Schulthess and Springer.

ABOUT THE AUTHOR

Rolf H. Weber is ordinary professor for civil, commercial and European law at the University of Zurich, Switzerland, and a visiting professor at the University of Hong Kong in China. His main fields of research are Internet and information technology law, international business law, media law, competition law and international finance law. He is director of the European Law Institute and the Center for Information and Communication Law at the University of Zurich. Since 2008, Rolf has been a member of the Steering Committee of the Global Internet Governance Academic Network and of the European Dialogue on Internet Governance. Since 2009, he has been a member of the High-level Panel of Advisers of the Global Alliance for Information and Communication Technologies and Development. He is an attorney-at-law, and his publication list is available at www.rwi.uzh.ch/lehreforschung/alphabetisch/weberr/person.html.

CHAPTER SEVEN: A PRIMER ON GLOBALLY HARMONIZING INTERNET JURISDICTION AND REGULATIONS

Michael Chertoff and Paul Rosenzweig

Copyright © 2015 by Michael Chertoff and Paul Rosenzweig

INTRODUCTION

We stand on the cusp of a defining moment for the Internet. Existing trends, left unaddressed, might very well lead to the legal fracturing of the World Wide Web. This brief chapter offers some thoughts on how this challenge should be resolved, concluding that multilateral agreement on a choice-of-law framework is essential to the continuing growth of the network.¹

THE PROBLEM DEFINED

The Internet is a globe-spanning domain. As of late 2014, more than 2.7 billion citizens of the world are connected to the network. Estimates vary, but somewhere around 500 billion different devices are also connected — and those numbers will only grow exponentially in the coming years.

The result is an increasingly common phenomenon: disputes and transactions that cross national boundaries. To be sure, the phenomenon is not new. There have been transnational commercial transactions (and transnational criminal activity) since the time that borders between nations were first created. But the growth of a system of near-instantaneous global communication and interaction has democratized the phenomenon of cross-border commerce in a transformative way that challenges and disrupts settled conventions.

The effect is most noticeable when we consider the intersection between private commercial activities and sovereign nations. Nations, quite naturally, seek to affect behaviour through laws and regulations that apply to individuals and corporations within their jurisdiction. But the growth in cross-border commerce is rendering traditional choice-of-law rules problematic at best. If one adds in the distributed structure of the network, inherent in the growing use of cloud architecture, the application of diverse legal systems to a unitary network becomes especially difficult.

For example, if a Swede stores his data with an American company that has a data centre in Canada, which country's law controls access to the data? What if (as is often the case given cloud architecture) his data is stored in more than one data centre located in more than one jurisdiction? When that same Swede provides personal information to a Chinese company, which then reuses the data for its own commercial purposes, where does he go to complain? And how is any actor to respond to inconsistent rules — where,

¹ We acknowledge at the outset of this chapter that such a framework may be difficult (some might say impossible) to achieve. We do not necessarily disagree. As outlined here, however, it is clear that the current situation, in the absence of such a framework, is untenable in the long run and destructive of economic prosperity. It may, indeed, be incapable of international resolution, but at a minimum, it is worth recognizing the necessity for action.

say, one country requires disclosure of data in a context that another country prohibits?

These jurisdictional problems are, if anything, confounded by the overtly political nature of many contemporary data disputes. They arise against a backdrop of authoritarian nations that want to control the content of the network and public concern about the extent to which nations undertake espionage in the cyber domain. These confounding factors are not strictly germane to the question of choice of law. Espionage is always illegal in the country in which it occurs, and content regulation is often more about political control than legal rules. But it would ignore reality to fail to acknowledge the contemporary political dynamic.

What we see today, in response to this conundrum, is an increasing effort by sovereign nations to unilaterally assert jurisdiction and control over matters they think are within their sphere of influence. These efforts fit under the general rubric of data localization requirements — the idea that data about, say, Germans must be stored in Germany and subject to German law. Even worse, such efforts are often ineffectual; although a nation such as Germany can demand localization, other nations are not obliged to honour that determination, and many nations (for example, the United Kingdom's Data Retention and Investigatory Powers Act) apply their laws extraterritorially.

A BRIEF NOTE ON JURISDICTION

At the heart of this problem lies the fundamental idea of jurisdiction: the question of which nation and which nation's laws may control the disposition of a matter. It reflects both a narrow power — that of a court to adjudicate a case and issue an order — and a broader concept of defining the territorial and lawful bounds within which a court, agency or government may properly exercise its power and authority.

Jurisdictional rules, of course, vary widely around the globe. There are often disputes about the legitimacy, in some broader sense, of a sovereign's assertion of jurisdictional authority.² Jurisdiction, in either sense of the word, is principally tied to the location of a person (including juridical persons, such as corporations) or things, as well as the subject matter authority to deal with an issue or dispute. Most nations have both courts of general jurisdiction that may hear any matter and courts that are limited to specific areas of subject matter expertise.

And so, when one characterizes the problem as one of jurisdiction, one is really speaking of power. Under what

² For example, in the United States (the jurisdiction with which the authors are most familiar), foreigners may be obliged to answer complaints in American courts only if they have a certain irreducible minimum of contacts with the jurisdiction such that they could reasonably anticipate the possibility of being called to account in that place. See *International Shoe Co. v. Washington*, 326 U.S. 310 (1945).

circumstances may the authority in one nation demand a response to its own legitimate inquiries? Given the complexity of the network and the increasing globalization of data transfers, problems of jurisdiction are multiplying rapidly.

TOWARD A SOLUTION

The current situation is untenable. Data localization and sovereign unilateralism will come with significant costs — both economic and social ones (Hill 2014). Global companies will be subject to competing and inconsistent legal demands, with the inevitable result that consumers will suffer diminished access to the network overall. Among other things, decisions about the location of servers and hardware will be driven by legal gamesmanship rather than technological or infrastructure considerations. The current free-for-all of competing nations needs to be replaced with an agreed-upon international system for a choice-of-law rule. What is needed is to harmonize existing rules within an agreed-upon framework of law.

What would such a framework look like? To answer this, it is useful to have a paradigmatic case in mind. Consider the case of an American company holding data about a European data subject at a data centre in Europe. When the US government seeks access to that data for law enforcement purposes, should the access be controlled by American or European law?³ What if they conflict?

One approach would carry the data localization movement to its logical conclusion and hold that the law of the country where the data resides controls access to it and rules relating to its processing. This parallels the usual case with physical evidence. Under such a system, for example, our paradigm case would be resolved by applying European law. This choice-of-law rule would have the virtue, at least, of clarity. Everyone concerned would know which jurisdiction's law would control.

But, in many ways, this clarity is illusory. In contemporary cloud structures, data is often stored in more than one location, either in disaggregated form or with copies resident in more than one data centre. It may also transit through multiple physical locations. A data localization choice-of-law rule would force corporations to alter the most economical structures of their data systems in order to secure legal certainty — an unnecessary cost. Alternatively, the data holders might choose not to take these costly steps, thereby creating the very legal uncertainty the rule is intended to avoid.

³ This paradigmatic case is modelled on a current dispute. See *In the Matter of a Warrant to Search a Certain E-Mail Account Controlled and Maintained by Microsoft Corporation*, M9-150/13 MJ 2814 (S.D.N.Y. July 31, 2014) directing Microsoft to disclose email maintained in Ireland. Other examples abound, including the move in Europe to require global compliance with EU privacy laws and the proposal in the United States to give court-approved search warrants global effect.

Perhaps more importantly, a data localization choice-of-law rule would create perverse incentives. Technologically, the most economically efficient place to store data is a product of a number of factors such as climate, infrastructure and proximity to users. With a localization choice-of-law rule we can anticipate at least two inefficient responses. First, some jurisdictions, either out of legitimate concern for their citizens or an authoritarian interest in control, will see this legal rule as a licence to mandate inefficient local storage requirements. Second, conversely, we might see other jurisdictions in a “race to the bottom” as they attempt to create data-access rules that are favourable to the data holders as a way of attracting business interests. Still others might develop rules that make them data-access “black holes,” where malicious actors can find a safe haven from legitimate scrutiny. None of these results is optimal, leading us to recommend against such a formulation of the choice-of-law rules.

Instead, we propose four alternate formulations that will also provide clarity in defining the jurisdiction that controls while being systematically less susceptible to economic gamesmanship and rent seeking than a data-location rule. We propose a choice-of-law rule based on either: the citizenship of the data creator; the citizenship of the data subject; one based on the location where the harm being investigated has taken place; or one based on the citizenship of the data holder or custodian.

A rule based on the citizenship of the data creator would tie jurisdiction over data to a familiar concept of personal jurisdiction — that is, the idea that one aspect of jurisdiction is the ability to exercise control over the person who created an item and, therefore, typically has ownership or control of the object or thing that is the subject of litigation. The overlap is not exact; sometimes the creator may not be the owner, in which case the interests of the creator may need to be distinguished, as they tend to be more personally direct than those of the owner.

In either instance, in most cases, citizenship brings with it universal personal jurisdiction — that is, the theoretical ability (often unexercised by a sovereign) to impose rules of conduct on a citizen wherever he may be in the world. The data creator or data owner citizenship rule would extend that paradigm in familiar ways such that those in control of data, wherever located, would be subject to the demands of their sovereign.

That rule may, however, be problematic, inasmuch as in the globalized economy the data creator or owner is often not the subject of the data. In other words, the data creator may be different than the individual whom the data concerns. For example, a photographer may be a data creator of a third party who is the data subject. In that case, the creator has an ownership interest, but it may be the latter who has the more compelling privacy interest.

Furthermore, an individual or corporate data owner with citizenship in one nation may store its data with a holder who has citizenship of another. Hence, the alternative of relying on the citizenship of the data holder would give primacy not to whose law the owner is subject to, but rather to the law of the entity holding the data — a result that will typically, but not always, apply the law of the physical server location where the data resides. This alternative would enhance geographic aspects of the network over legal ownership perspectives. In addition, it may be valuable to recognize that holders who disclaim ownership will likely be treated in a manner different from those holders who also take an ownership interest in the data that they hold.

A rule that focuses on the citizenship of the data subject would serve to elevate personal jurisdiction as it relates to the individual or corporation to whom the data most directly relates and who often, although not always, is the creator of the data in question. This alternative would serve to enhance personal control of data, at the expense of degrading the comparative value of a sovereign's control of the data owners or data holders subject to its jurisdiction.

Finally, jurisdictional rule that determines the result based on the locus of the harm would reflect a sea change in current trends, away from jurisdictional assertions based on status. It would, instead, substitute a predominant effects-like test of jurisdictional primacy that would be more flexible and indefinite, with uncertain application. It would, however, be more certain in addressing legal harms caused by the conduct that is the underlying subject of inquiry.

To be sure, all of these rules will have grey areas at the margins. Some data subjects may be dual citizens. Some data holders may have corporate headquarters in more than one nation. And some events may give rise to harm in more than one location. But none of these are circumstances that are as readily capable of manipulation as data location; indeed, in many instances they will be extrinsic to the data and the product of other circumstances.

There are sound arguments for and against each of these possible rules:

- A rule based on citizenship of the data subject would ground Internet jurisdiction in a familiar legal construct. It would focus on the individual, whose privacy rights and activities are likely to be most directly implicated by any jurisdictional rule. It would also reinforce the idea that citizenship and sovereignty are closely linked. It might, however, be the most difficult rule to implement technologically, since data often does not have a flag for citizenship of origin or ownership and retrospectively adding such a marker might prove challenging, if not impossible.
- A rule based on citizenship of the data holder would have the virtue of ease of application — a single rule would apply to all data held by the data holder. It would also, however, have the unfortunate effect of creating transnational incentives of the same sort as a data localization rule, with the added consequence of fostering economic nationalism. And data holders who are also data owners may have greater obligations than those who are not owners. Finally, since localization rules are not self-executing, their adoption may increase confrontation at the cost of cooperation and will ultimately have harmful effects on innovation and economic development.
- Much the same would be true of a rule tied to the citizenship of the data creator or owner. Such a rule would incentivize unilateralism at the expense of cooperation. Moreover, where the data owner and data subject are different, focusing on the former for jurisdictional purposes might have the unintended effect of undervaluing privacy values.
- A rule based on the location of harm seems the least capable, generally, of manipulation and most directly linked to cognizable sovereign interests. It suffers, however, from the ability of sovereigns to define and manipulate the definition of harm, and would only be implementable for certain universally agreed upon harmful acts, such as murder.

None of these rules is perfect. Each is capable of manipulation and each will require some transnational cooperation to implement. When the matter involves an inquiry outside the jurisdiction of the nation seeking the data, requests for assistance under each of these rules will have to be processed through cumbersome and possibly unavailing mutual legal assistance treaty (MLAT) channels. This means that some jurisdictions will continue to serve as safe havens for malicious actors no matter what choice-of-law rule is chosen.

Accordingly, concurrent with a revision to the choice-of-law rules, we would be wise to develop a more streamlined MLAT structure. If countries could rely upon the prompt response to data requests, they would be less inclined to act unilaterally in the assertion of jurisdiction. Better MLAT responsiveness (combined with reciprocity obligations) would minimize the temptation to create safe harbours through data localization. It would also lessen the adverse effects of a new rule on law enforcement that would result from adopting one of our possible jurisdictional approaches. MLAT reform would assure law enforcement that, in the end, despite jurisdictional rules that would limit its ability to act unilaterally, it could still avail itself of a reformed MLAT process for an effective response to criminality.

The virtue, however, of these suggested rules lies in their ability to create clarity and ease of use among willing participants. We could, for example, imagine a transnational agreement on data availability tied to the protection of life and property, perhaps with some degree of judicial oversight, which could be implemented throughout the West. That limited goal itself would be a major achievement in creating security, clarity and consistency on the network.

WORKS CITED

Hill, Jonah Force. 2014. "The Growth of Data Localization Post-Snowden: Analysis and Recommendations for U.S. Policy Makers and Industry Leaders." *Lawfare Research Paper Series 2* (3). www.lawfareblog.com/wp-content/uploads/2014/07/Lawfare-Research-Paper-Series-Vol2No3.pdf.

ABOUT THE AUTHORS

Michael Chertoff is executive chairman and co-founder of The Chertoff Group, a premier global advisory firm focused exclusively on the security and risk management sector. In this role, he provides high-level strategic counsel to corporate and government leaders on a broad range of security issues, from risk identification and prevention to preparedness, response and recovery. During 2004 to 2009, he served as secretary of the US Department of Homeland Security, where he led the federal government's efforts to protect the United States from a wide range of security threats. Earlier in his career, Michael served as a federal judge on the US Court of Appeals for the Third Circuit and head of the US Department of Justice's Criminal Division where he investigated and prosecuted cases of political corruption, organized crime, corporate fraud and terrorism — including the investigation of the September 11 terrorist attacks.

Paul Rosenzweig is a senior adviser to The Chertoff Group. He previously served as the deputy assistant secretary for policy and as acting assistant secretary for international affairs at the US Department of Homeland Security. During this time, Paul developed policy, strategic plans and global approaches to homeland security, ranging from immigration and border security to avian flu and international rules for data protection. He currently provides legal and strategic advice on cyber security, national security and privacy concerns to individuals, companies and governments.

**SECTION FOUR:
BALANCING TECHNICAL OPENNESS
AND FRAGMENTATION**

CHAPTER EIGHT: MARKET-DRIVEN CHALLENGES TO OPEN INTERNET STANDARDS

Patrik Fältström

Copyright © 2016 by Patrik Fältström

ACRONYMS

AIN	Advanced Intelligent Network
API	Application Programming Interface
CPE	customer premises equipment
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
HTTP	Hypertext Transfer Protocol
IEEE	Institute for Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IMAP	Internet Message Access Protocol
IoT	Internet of Things
IP	Internet Protocol
IRC	Internet Relay Chat
ISOC	Internet Society
ISPs	Internet Service Providers
ITU	International Telecommunication Union
ITU-T	ITU Telecommunication Standardization Sector
MIME	Multipurpose Internet Mail Extensions
NAT	network address translation
POP	Post Office Protocol
RFC	Request for Comments
SMTP	Simple Mail Transfer Protocol
TCP	Transmission Control Protocol
TLD	top-level domain names
W3C	World Wide Web Consortium
XMPP	eXtensible Messaging and Presence Protocol

INTRODUCTION

The ideal of an “open” Internet is often portrayed as wishful thinking — something that would be nice to have in a perfect world, but is not always compatible with the need for revenue and the harsh reality of the market economy. This is a profoundly false impression, because the openness of the Internet and its mechanisms, from the development of technical standards to the operation of the global network, confers enormous practical economic benefits.

In practice, the open Internet has been fertile ground for the invention and development of remarkable new companies, capabilities and modes of human interaction. The openness principle continues to guide the Internet’s evolution in technical, economic, political and social dimensions. Innovation in the open Internet is achieved by

consensus through open collaboration among researchers, manufacturers, service providers and users. Innovation can start anywhere and propagate in any direction.

But that’s the long-term view. In the short term, market forces can drive fragmentation and anticompetitive “silo” approaches to product and standards development that erode the foundations of the open Internet. These forces are not only short term with respect to commercial advantage but also short-sighted regarding sustainable innovation and economic growth. They can be countered by a clear understanding of the tangible benefits of the Internet’s traditional open approach to standards development.

INTERNET FUNDAMENTALS

The Internet, like other communication networks from highways to telephones, consists of a web of connections, and anyone who is connected to the Internet can communicate with anyone else. The basic communication is via Internet Protocol (IP) packets, each of which carries a small amount of data; many of them together carry a document, a movie, a telephone call, a recipe or the latest photographs of someone’s cat.

The point (or “node”) at which each device is connected to the Internet is identified by an IP address, and because these addresses are globally unique, communication between two devices is, in its simplest form, a sequence of IP packets sent back and forth between them, in which each packet includes the IP addresses of both. The service we get with Internet access is therefore relatively simple: the best-effort delivery of IP packets from the source to whatever node has the address stated as the destination address in the packet.

The interpretation of the contents (payload) of each packet is up to the software in the nodes that send and receive the packets. This is why, when comparing traditional telecommunication and the Internet architecture, one says that the intelligence — the knowledge of what services exist — has moved from the network to the edge. Just by changing the software in two nodes that communicate (and without changing anything in the network), a new service can be launched. We call this “permissionless communication and innovation.” Innovation and launch of new services has moved from being a business for the owner of the network to a business for whomever controls the devices connected to the Internet — the end user.

END-TO-END COMMUNICATION

The end-to-end communication between two Internet users is often referred to as “peer-to-peer,” in part to distinguish it from “client-server” communication between an end user (the client) and a broadly available application service provided by someone else, such as Google, Netflix or Amazon (the server). From an Internet architecture

point of view, these service companies also have nodes with unique IP addresses with which their customers (the end users) communicate. Both servers and their clients are connected to the Internet by companies that provide the infrastructure to move IP packets from one node to another. It is very important to distinguish the packet-level exchanges facilitated by these Internet service providers (ISPs) from the application-level services, such as “movie watching on demand,” that are offered by content provider companies such as Netflix. In this example, both Netflix and its clients are separate customers of ISPs.

GLOBAL UNIQUENESS

The ability to exchange IP packets unambiguously between any two nodes on the Internet requires that all nodes have globally unique IP addresses. To make path selection easier, IP addresses are generally allocated according to network topology, so that two nodes that have addresses adjacent to each other are located within the same network. Because of this, IP addresses might change when a node is moving in the network. Domain names create longer-term stability in addressing, and make the addressing more human friendly. Each domain name is also globally unique, and the distributed database called the Domain Name System (DNS) maps domain names to IP addresses.

We Have One – and Only One – Set of Domain Names

Domain names are sequences of LDH-string¹ labels organized into a tree-structured hierarchy in which the manager of a domain at one level can allocate subdomains at the next level, moving away or “down” from the root of the tree. At the root, we guarantee the global uniqueness of all domain names by having one and only one manager of the top-level domain names (TLDs) (Request for Comments [RFC] 2826).² Because each domain name is unique, when it is used to address a node on the Internet, it will lead unambiguously to one and only one node (or to no node at all, of course, if no node with that name exists). A given domain name, assigned to a domain holder, is therefore recognized uniquely all over the world.

We Have One – and Only One – Set of IP Addresses

IP addresses are also assigned uniquely and unambiguously to Internet nodes around the world. A given IP address will lead to one and only one node (or to none). Each IP address is assigned through a system of IP

address registries (RFC 7020)³ to just one party, and is via announcement⁴ recognized uniquely all over the world.

OPEN STANDARDS

Having one and only one set of domain names and IP addresses enables any Internet-connected entity to identify and communicate with any other entity connected to the Internet.

But communication relies not only on the ability to convey information from one party to another; the parties must also understand each other. Both the syntax (the data format) and the semantics (meaning) of the information exchanged must be understood. The information consists of *commands* (or requests for action) and *data* (to which actions are applied). Standards ensure that all the parties interpret the commands and data in the same way (Bertin, Crespi and Magedanz 2013).

Traditional Telecommunications

Traditional telecommunication networks involved a central “master” system with which multiple “client” systems communicated directly. The master dictated how the clients communicated. If two clients wanted to communicate with each other, they did so by first connecting individually to the master. Examples of such networks include traditional telephony and banking systems.

The standard specifying how clients could communicate was therefore defined for these networks in terms of how to communicate with the central master. Often this would be expressed in the form of an API, typically a software library designed to be incorporated into each client system.

In such a master-client arrangement, the connection from any client to the central master was cheap (in both resources and cost), but the duration-based charging scheme ensured that the total transaction cost to clients was high. A simple example is that lifting a telephone receiver and receiving a dial tone was cheap, but the cost for an actual call to another telephone was high, based on payment by the minute as the call proceeded.

Quite often, service providers developed both the master system and part or all of the client systems, and how the actual communication took place was a proprietary (commercial) secret. Providers competed with each other on both price and the efficiency of the protocol used to communicate with clients. This of course included optimizing the protocol so that production and operating

1 The acronym “LDH” stands for “letter digit hyphen,” the three types of characters that can be used to form a syntactically correct domain name label. The use of other character types has been introduced through the “Internationalized Domain Name” program, but is beyond the scope of this chapter.

2 See <https://tools.ietf.org/html/rfc2826>.

3 See “The Internet Numbers Registry System” (August 2013), <https://tools.ietf.org/html/rfc7020>.

4 “Announcement” refers to the way in which the routing protocols of the Internet distribute knowledge of where each of its nodes is located, so that IP packets can be sent from one router to the next on a path that eventually ends at the correct node.

costs were as low as possible. In this way, even a nominally “standard” protocol endured many proprietary modifications as providers sought competitive advantage.⁵

In such an environment, a provider could attract a competitor’s customers by implementing that competitor’s API (perhaps by reverse engineering or licensing), enabling communication between its own clients and those of its competitor even though they did not use the same protocol for communication. In this case the provider’s master acted as a proxy between its protocol and the one used by its competitor. The same effect could of course also be achieved by an agreement between the two providers to exchange traffic between their master systems.

Internet Communication

In the Internet, end nodes can communicate directly with each other without the intervention of a central master — the network intelligence is implemented in the nodes themselves. This requires the end nodes to implement the same protocol, and for the Internet this was ensured by explicit and uniform standardization of the communication protocols in terms of “the bits on the wire.” Any node that correctly implements the standards can communicate with any other node that does the same.

The difference between specifying a protocol for node-to-node communication and one for node-to-master-to-node communication may seem small — the end result is the same — but if one looks at how the specifications are developed, there is a big difference. When the protocol between nodes is specified, the development of the standard is likely to take place in an open environment in which many parties are involved. When the specification of how to communicate with a central master is developed, it is often controlled or determined by whoever owns and operates the master(s), and “clients” have little influence over the standards they are then forced to adopt.

As the Internet model of end-to-end communication has displaced the centrally controlled master-slave configurations of traditional telecommunications, the process of developing and using a protocol standard has become more *open*, in everything from participation in the creation of the standard, to access to the standard itself, to licences needed for implementation. This is no accident — an open standards process produces results that incorporate a broad range of ideas and perspectives (not just those of a single central authority); it can be implemented and used by anyone (not just those who buy

into a proprietary scheme that works only within centrally controlled boundaries); and it establishes a level playing field on which competition is not distorted by proprietary advantage. Internet standards are open not because some authority decided that they should be, but because an open process produces standards that are better — technically, economically and politically — for all of the participants (ISOC 2015a).

OPEN STANDARDS DEVELOPMENT

The Internet Engineering Task Force (IETF) is often presented as a good example of a standards body that fulfills the requirements and expectations for an open standards process.⁶ In 2012, along with the Institute for Electrical and Electronics Engineers (IEEE), Internet Society (ISOC), World Wide Web Consortium (W3C), and Internet Architecture Board, the IETF endorsed the “OpenStand Principles”⁷ of due process, broad consensus, transparency, balance and openness. These principles codify the six key features or “abilities” that characterize and define an open standards process:

Ability to Participate in Development of the Standard

The ability to take part in the development of a standard has two aspects: whether only certain parties can participate, and if the cost of participation is high. In many traditional telecommunication standards development processes the cost of participation is high and there is no ability to participate if you are the wrong kind of entity (regardless of how much you pay).

Ability to Access Working Documents

Even if direct participation is not possible, a standards development process might arrange for “outsiders” to review preliminary documents — perhaps because the direct participants in the process want input from others. Non-members might be interested in early drafts of a standard so that they can make earlier decisions on whether to implement the standard and assess how much the standard may affect their business. Some standards organizations do give access to all documents while others do not. For example, in the Telecommunication Standardization Sector of the International Telecommunications Union (ITU-T), some members (including Sweden) have argued that all documents should be freely available; however, a majority of ITU-T members object to free access.

⁵ “To summarize, advanced intelligent network (AIN) equipment supplied by different vendors will normally not work well together. Although global network operators cannot derive any competitive advantage from different network systems of similar price and quality, they can derive such advantage from integrating these into more seamless structures and services (bundled, customized, and mass-produced) (Bohlin 1998, 109).”

⁶ See www.ietf.org/about/process-docs.html.

⁷ See <https://open-stand.org/about-us/principles/>.

Ability to Participate in Decision Making

Even where participation in a standards process is allowed (including access to working documents), it is sometimes hierarchical in that only certain membership types can participate in formal decisions (such as whether to approve a particular draft as a standard), which often are made by vote. This gives some members greater power than others to influence the final outcome of the development of a standard.

Ability to Appeal

If there is an error in the specification of a standard, or if there is a view that the specification does not solve the problem it was supposed to solve, it is essential that the process provide for appealing the decision to approve it. An appeal can lead to a change in the developed standard or to initiation of development of a new standard that replaces the old one. It can of course also be rejected.

Ability to Access the Standard

After a standard has been approved, it must be accessible to those outside of the standards development process so that it can be implemented. The business model of some standards bodies relies on control over how the standard is distributed and how much access to the standard should cost (perhaps graded according to the type of access or the type of entity seeking access). The product of an open standards process, however, must be freely available to all.

Ability to Implement the Standard

Even where access to a standard is freely available, some or all of the standard might be encumbered by intellectual property rights (such as patents) that must be licensed from their owner(s). In some cases, the licensing fee is defined by the standards organization (often an industry consortium); in other cases, it might be defined by the owner. For open standards, it is customary for the rights holder to grant an implied licence to anyone who wants to implement the standard.

STANDARDS EVOLUTION

Two examples of Internet applications for which open standards have been developed — electronic mail (email) and social media instant messaging, or “chat” — illustrate how standards evolve.

ELECTRONIC MAIL

Electronic mail, with its familiar “user@example.com” addresses, is probably the most widely used and recognizable Internet service. In Internet terms, its

protocols are ancient — the first email standard⁸ was published in 1980. It has been updated regularly since then, maintaining backward compatibility with previous versions at each step.

Email Standards

The basic standard for email consists of two core specifications: one that specifies how an email message is formatted (with To, From and Subject headers, for example), and one that specifies how to send an email message to a recipient. The first standard for the Simple Mail Transfer Protocol (SMTP) was published in 1982.⁹ It specifies a protocol in which a client opens a Transmission Control Protocol (TCP) connection to a server, interacts with the server via commands and responses and then closes the connection. Later, as email evolution led to configurations in which clients were not directly connected to servers, two more protocols — the Internet Message Access Protocol (IMAP)¹⁰ and the Post Office Protocol (POP)¹¹ were developed to facilitate email retrieval.

But after evolving from direct interaction with the message store to using POP and IMAP, email clients have evolved further to become “webmail” clients, which use the World Wide Web’s Hypertext Transfer Protocol (HTTP) to manage email interactions. One such client consists of a web browser that accesses a normal web page that is dynamically created by some software on the server side of the HTTP connection. This software, not run by the client, often in turn acts as an IMAP client.

Because of this, email has evolved back to a system in which the email user is connected to some application using a mechanism that is very similar to the old direct message store connections, although the connections are now made over the Web using a web client.

Standard Email Extensions

As email use expanded, the standards evolved to accommodate new demands to support non-ASCII text, images and other data formats, either as “attachments” to or in the main body of email messages. Beginning in 1991, the IETF developed a highly flexible standard for Multipurpose Internet Mail Extensions (MIME)¹² that provided support for text in character sets other than ASCII; non-text attachments such as files containing audio, video, still images and applications; and messages

8 See RFC 772, <https://tools.ietf.org/html/rfc772>.

9 See RFC 821, <https://tools.ietf.org/html/rfc821>. The most recent full specification of SMTP is RFC 5321, <https://tools.ietf.org/html/rfc5321>.

10 See RFC 3501, <https://tools.ietf.org/html/rfc3501>.

11 See RFC 1939, <https://tools.ietf.org/html/rfc1939>.

12 See RFC 2045, <https://tools.ietf.org/html/rfc2045>.

with multiple parts (so that a single message could include multiple text and non-text elements). It also supported privately defined extensions, so that if someone wanted to send data structured in a way known only to them, they could do so by tagging the data with a private name.¹³

MIME could have led to an explosion of private email structures, but it has not. Instead, people are trying to use a small set of common formats wherever possible: one for still images, another for video, a third for signatures and so on.

The high degree of interoperability that has been achieved by the standardization of SMTP, POP and IMAP has led to a rich marketplace of server and client software, including “webmail” that behaves as if it were an IMAP client. MIME has further enabled an extensions mechanism whereby extensions can be either standardized and interoperable or non-standardized and proprietary. This has led to a situation in which an implementer can choose from a wide variety of interoperable and proprietary email configurations.

Non-standard Email Exceptions

As we might expect of such a widely used and economically significant system, email has not evolved uniformly in the direction of interoperability. Clients today are faced with choices that go beyond the standardized IMAP and POP — they may instead choose, for example, “Exchange” or “Google.” Microsoft and Google are not the only players pushing email in proprietary directions, but they serve as useful examples of the way in which market forces affect the evolution of a standard.

Microsoft Exchange

The Exchange server (and the client Outlook) can use IMAP and SMTP for communication, but for full functionality they also implement a proprietary protocol between clients and servers, which includes specific authentication mechanisms. Whether the standardized protocols should be enabled or not in the Exchange server is a configuration option, and many system administrators turn these options off. That way, if someone uses an email service that in turn is implemented with the help of Exchange, they must use the Outlook Client. It is also the case that if they want to use the Outlook Client with full functionality, they must use the Exchange server.

Google Mail

Google implements their email service by exposing both a web interface and an IMAP interface to clients. However, they use IMAP in an innovative way: to categorize mail

by using tags, and exposing that to the client as folders (or containers). By tagging an email message with more than one tag, it can appear in more than one container. For this user experience to be fully realized, the client must understand Google’s extension to IMAP, and many IMAP clients do indeed include support for this; however, it is not a standard IMAP way of managing tags and folders.

SOCIAL MEDIA

The term “social media” refers to a wide range of applications and services. In this section we are interested only in the instant message, or chat, feature of most social media platforms.

Internet Relay Chat

Long ago in the time frame of the Internet — in the early 1990s — a text-only instant messaging system called Internet Relay Chat (IRC) was invented in Finland. Anyone could set up an IRC server. These service providers “peered” with each other, and the addressing was based on the name of the service plus the name of whatever was to be addressed — an individual or a chat room (or “channel”). IRC was popular and is still used by some programmers. The protocol is simple, and it is very easy to create robots that respond to messages in the various channels that in many cases act as permanent chat rooms.

IRC was defined not by a standard but by its implementation in open-source software. Anyone could look at the source code and develop either server or client software. To explain how the protocol works, an experimental RFC¹⁴ was created. But IRC still evolves as a constellation of mostly open-source implementations. Informational RFCs are released now and then explaining updates to the protocol, but they are not uniformly adopted. New features arise as enhancements to an implementation that “catch on” with other software developers, some of them coordinated by more formal groups such as the IRCv3 Working Group.

Jabber

Interest in a formally standardized chat protocol developed in the late 1990s, when the IETF launched a working group to develop one. No consensus could be reached in this working group, so instead of a single standard it published several of the contenders as experimental RFCs — the idea being to allow the community to implement the protocols, and find out from experience which one should win. Which, in fact, none of these did.

Instead, an instant messaging system called Jabber was developed outside of the IETF in 1999, as an open-

¹³ To avoid confusion among privately defined mail extensions, the IETF defined a registry and registration procedures for message header fields in RFC 3864, <https://tools.ietf.org/html/rfc3864>.

¹⁴ See RFC 1459, <https://tools.ietf.org/html/rfc1459>.

source software implementation.¹⁵ The Jabber developer community specified an eXtensible Messaging and Presence Protocol (XMPP),¹⁶ which in 2002 was moved to, and accepted by, the IETF. Jabber began very much the way IRC did, but followed a different route; today XMPP is the subject of a full set of Internet standards overseen by an IETF working group, and is the dominant interoperable chat protocol.

Proprietary Protocols

More recently, the instant messaging and related services launched by Facebook, Twitter, Skype and Google (for example) have been based on proprietary rather than standard protocols. No description of the protocol used among the various parties that communicate is provided, thus there is no ability for a third party to implement (or even interact with) the service. Access to the Facebook service, for example, is available only from Facebook.

This model differs dramatically from the provisioning of services based on standards. To get email service or instant messaging, we can choose from a multitude of providers, each of which in turn can choose from a multitude of different providers of software — or write their own. The difference between the open development and evolution of IRC and Jabber and the current growing reliance on entirely proprietary alternatives has enormous consequences for Internet users. In a world of proprietary, non-interoperable services, users are limited to choosing either Facebook (for example) or Google — they cannot choose among alternative providers of “Facebook service” or “Google service.”

MARKET FORCES

As we look at how standards have evolved, we see that the developers of software and services have cooperated in producing open standards that have led to interoperability. This has created a competitive landscape in which no single player can completely dominate the market. Thousands, if not millions, of providers of web hosting have appeared, for example, and the range of available server and client software for open-standard applications is extensive.

But providers of server software have always also had an economic interest in controlling clients, and the business models of large service providers have always favoured anti-competitive domination of a market over competition (GECON 2007). Absent an alternative countervailing value proposition based on economic advantages to these and other businesses, market forces will drive the evolution of

Internet protocols and services away from interoperability and toward user “lock in.”

Fortunately for users, research suggests that such value propositions in favour of openness do exist (Zhu and Zhou 2011). Two modern developments — the Internet of Things (IoT) and the “cloud” — illustrate how market forces traditionally operate against open interoperability and how they can be redirected.

THE IOT

As the latest hot topic in the technology industry, the IoT has produced a mountain of commentary and analysis that runs the gamut from the breathless excitement of optimists to the dark warnings of pessimists.¹⁷ From the standpoint of the Internet architecture, the IoT is hardly a new idea — it is simply devices connected to the Internet, each with its own IP address, just like always — but from the standpoint of our assumptions about how the Internet is used, it is indeed a radical departure.

Although it has always been “things” that are actually connected physically to the Internet, most models of interaction have taken for granted that at least one of the parties to any Internet communication is a person — a “user.” In the traditional Internet, communication may be user to user, or user (client) to computer (server). The IoT adds the third option of computers talking to each other without human intervention. And this third option may involve much more than talk — after all, completely autonomous sensor networks have been gathering and storing data without human intervention for decades. What is new in the IoT is that connected devices may also autonomously analyze the information they exchange and take actions independently as a result.

The emergence of the IoT owes more to companies’ marketing incentive to make devices ever more functionally “intelligent” than to any collective sense within the Internet community that things should be able to talk to one another. What standards exist, therefore, tend to be developed by individual companies or industry consortia that focus on enhancing the capability (and therefore marketability) of the “thing” without particular regard to interoperability with “things” in other industry sectors — or with the “things” of competitors.

Case Study: Lighting Control

A simple example of an arena in which open standards are missing is control of light bulbs. Superficially, it is a simple problem to control a light bulb — the traditional method uses two wires, one live and one neutral, and a switch turns power on or off on the live wire. In modern lighting

15 For a history of Jabber and XMPP development compiled by the XMPP Standards Foundation, see <https://xmpp.org/about/history.html>, including the involvement of the IETF’s XMPP Working Group, see <https://tools.ietf.org/wg/xmpp/charters>.

16 See <https://tools.ietf.org/html/rfc6120>.

17 For a summary of the “promise” and “peril” scenarios see Delic (2016).

systems we can often also control the intensity and colour of the light; this can be done by using an overlay protocol on the existing wires, of course, but it is even easier to do if the light bulb is connected to a network and has a unique address, at which controlling software (acting as the light switch) can communicate with it (and perhaps many other devices) using a standard protocol.

This still sounds simple, but the problem is that there is no such standard protocol for talking to light bulbs, and no system of unique addresses for them — so no light bulb and switch interoperability, and to make matters worse, the light bulbs themselves are not interchangeable. A closer look at two examples will make this problem clearer.

Philips Hue System

In the Philips Hue system¹⁸ the light bulb communicates with a gateway (the ZigBee Bridge) using a proprietary protocol over ZigBee,¹⁹ a low-power digital radio technology based on the IEEE 802.15.4 standard for wireless personal area networks.²⁰ A light-control application communicates with the same gateway using a different proprietary IP-based protocol that is not documented by Philips, although third parties have reverse-engineered the protocol and developed libraries for a variety of programming languages (including perl, php, and python) that can be used by application developers.

A light switch must understand at least one of these two proprietary protocols — the one that runs on top of ZigBee to communicate with light bulbs, or the one that uses IP to communicate with a Hue gateway. If Philips changes the protocol, the light switch has to be updated. And, of course, unless the light switch update takes place at the same time as the Philips protocol change, there will be some interval during which the switch can't control the light. Although neither Philips nor the light switch manufacturer wants this to happen, there is no well-defined change control for the Philips protocols that includes third-party suppliers of light bulbs, switches or control applications.

LifX System

The LifX²¹ system uses standard WiFi, rather than ZigBee, and runs IP directly from one device to another without an intermediate gateway. In LifX configurations the devices — light bulbs and switches, and also many other devices using the “If This Then That” web service — connect to the local wireless network and get IP addresses using Dynamic

Host Configuration Protocol (DHCP). The protocol used, including the encryption, is defined by the manufacturer of the light bulb and is not publicly available. Some reverse engineering has been done to provide alternatives, but most popular access to the light bulb is via the application developed by LifX itself.

IoT Standards

The pressure for manufacturers to build “silos” — vertically integrated families of devices that talk to each other, but not to devices made by other manufacturers — is evident in this case study. Lighting control is one of the simplest and most common examples of an IoT application, and because it is a consumer-oriented technology, we would expect it to be based on standards that create interoperability, at least at the level of the simple devices (bulbs and switches) that are mass-marketed to the public. But each company imagines that its proprietary approach will become widely adopted as the “de facto” standard, with respect to which it will have an obvious competitive advantage over other companies pursuing the same “maybe it will be me” strategy. Interoperability and openness are actively detrimental to such a strategy, because they dilute the advantage that a company expects to have when “everyone” starts using its version. Consumer electronics has evolved in this way for many decades; there is no reason to expect that IoT evolution will take a different course (Blind 2004).

Only by using open standards can the light bulbs and the controlling software be made interoperable, enabling competition that could foster innovation and evolution. Today, the lack of interoperability has severely limited the growth of IP-based connected light bulbs.

THE CLOUD

The term “cloud computing” refers to a shared-resource model in which individual computing devices obtain application, platform and infrastructure services such as computation, storage and software via network access to a server — or, more commonly, a distributed “cloud” of servers that collectively provide those services. In the context of this chapter, we are interested in a particular feature of cloud computing: the way in which it can serve as an intermediary, or proxy, to relay communication between devices that are connected to the Internet in a way that prevents them from communicating directly with each other.

Network Address Translation

The Internet's system of global addressing supports — in principle — the end-to-end connectivity of any two Internet-connected devices. In practice, however, the most common connectivity arrangement for residential and business premises has one device — in telecom

18 See www2.meethue.com/en-us/about-hue/what-hue-is.

19 See www.zigbee.org/what-is-zigbee/.

20 Although it is not a formal standards body, the ZigBee Alliance (see www.zigbee.org/zigbeealliance/) is the focal point for most ZigBee technology development.

21 See www.lifx.com/.

terminology, the customer premises equipment (CPE) — actually connected to an Internet access provider and all other devices at those premises connected through the CPE using network address translation (NAT). The CPE typically receives one IP address via DHCP from the access provider, and shares it with all the other devices, which do not get individual IP addresses of their own. The CPE’s IP address is dynamically allocated by the service provider, so it is not associated with an Internet domain name in the DNS. And the CPE also typically acts as a firewall, filtering traffic so that all communication with its attached devices must be initiated by them.

The consequence of this arrangement is that the devices connected through such a CPE cannot actually be reached directly from the Internet, and only the CPE, with its dynamically allocated IP address, can be reached from outside. All Internet communication must therefore be initiated by the devices themselves; they communicate directly with the CPE, which uses its own IP address and a local identifier to set up the path to the other Internet-connected device and manage the subsequent communication between them.

End-to-Cloud-to-End

In such a NAT configuration, the only way that information can be exchanged between two devices is if each device opens a connection through its CPE to a server that manages the flow of data between them. Two devices configured with NAT cannot communicate directly using IP.

All “end-to-end” communication is therefore actually store-and-forward, with storage in “the cloud” as an intermediary. As cloud storage initially was created to solve the ability to communicate (or lack thereof), the specification of the protocol used does not have to be published; the cloud service is created by the same party that created the device. The communication is internal to the service, and no global communication exists. No protocol standard is needed for the (non-existing) communication.

Application Programming Interfaces

In a NAT environment, user-to-user communication is mediated by a centralized service “in the cloud.” The service itself defines how to interact with it by specifying an Application Programming Interface (API). This specification tells devices how to use the service, which is a very different thing from a protocol standard that specifies the way in which two users may communicate end-to-end. As the licences, terms and conditions associated with these APIs are defined by the provider of the service, the end users have little choice.

This can be viewed as a classic example of a one-sided market. For example, the service provider can change the API at any time. In practice, this will always come as a surprise to its customers, whether or not its contractual agreement with the API user says that changes will be announced before being made, or that those announcements actually are made.

Data Collection

A significant market force driving the interest in service silos defined by APIs rather than end-to-end protocols is the value of what has come to be called “big data” — collections of enormous size that have only recently become susceptible to analysis (Chen et al. 2014). With the advent of tools that make feasible calculations on entire very large data sets (rather than on smaller statistical samples), being a proxy through which communication between end users takes place has become valuable. Today we see companies just collecting data, even if they do not know what calculations to make (yet); the data sets have become valuable in themselves, creating a revenue opportunity for the service provider that in some cases can compete with the sales of the service itself.

Collecting and selling the data can also allow a service provider to lower or eliminate the fees it charges to use the service. This is naturally popular with consumers, who today in many cases enjoy the use of cloud-based services for free. But the easily recognized advantages of “free” make it harder to engage the more difficult issues of data “ownership,” including access, privacy and sharing data with third parties.

AN OPEN INTERNET FUTURE

This chapter has presented examples of the way in which market forces can lead to fragmentation of the nominally global and open Internet into service-oriented silos. In this concluding section we argue that the silo scenario can be avoided, and that the values of an open Internet can be extended into the future by recognizing and promoting forces that counter market forces.

THE CHALLENGE

If technical constraints (such as the Internet Protocol version 4 address length limit that led to the widespread deployment of NAT) make end-to-end communication too difficult, then users will turn to proxies that involve intermediaries in the end-to-end path. User-to-user communication via proxy introduces opportunities for third-party control, access to content and metadata, and charging. If on top of this the protocol is proprietary, then all devices must communicate with the same central cloud service. The proxy provider is in full control. From a business standpoint, of course, this sort of control is extremely valuable, and many companies today are

competing vigorously to become the preferred proxy for the household.

The best-case scenario in such a third-party dominated configuration would be that devices from different manufacturers are able to communicate with and via the same proxy. But even in this case, multiple proxies may provide services for the same household. The lack of standard protocols, both between devices and between devices and cloud services, leads to the implementation of services as isolated silos. Even the APIs that define the silo services will exist for only as long as the corresponding cloud services exist. In practice, this also limits the lifetime of the devices that are sold to connect to the service, as the device itself might still be functional even if the service is turned off.

RECOMMENDATIONS

The protocols that have been developed within the Internet architecture are deliberately peer-to-peer. Even those that specify client-server interactions, such as the email protocols POP and IMAP, specify interactions at one level without constraining the way in which other parts of the system may be defined or implemented. Silo services define only an API that governs the entire spectrum of interaction with users. The most important recommendation for avoiding a fragmented Internet future is to promote the deployment of communication systems based on standard protocols rather than service-specific APIs.

The most broadly useful and valuable protocols are those developed by open standards processes in which everyone can participate and to which everyone can contribute. Protocols that depend on privately owned intellectual property may be subject to a variety of different licensing terms, but as with the protocols themselves, the more open the licensing terms, the more beneficial the results in the market. APIs that are specified as part of a peer-oriented (as opposed to silo-oriented) system should also be developed by an open standards process.

The gold standard for an open and transparent standards process has been set by independent organizations such as the IETF, the W3C and the IEEE, but industry alliances such as the Internet Protocol for the networking of Smart Objects Alliance or the Industrial Internet Consortium can also develop open standards. Industry-sponsored standards efforts do not always welcome the participation or contribution of the users who will be affected by their outcome, but industry leader collaboration is likely to at least minimize the number of silos and increase device interoperability for the end user.

CONCLUSION

Public sector organizations should use every opportunity that arises in procurement, regulation and project funding to require the use of open standards when they are available and to promote their development when they are not. This responsibility is especially important for socially critical systems such as electronic identification and payment schemes, for which the third-party control feature of service silos is unacceptable.

The market forces that favour service-oriented vertical integration over a disintermediated open Internet create strong economic incentives for individual companies to build silos with APIs rather than interoperable devices that implement standard protocols. Countering those forces to preserve the broad economic and social benefits of an open Internet for its users will require awareness and effort on the part of users and their public sector organizations, and a willingness to take a longer view of their business interests on the part of individual companies and industry consortia.

WORKS CITED

- Bertin, Emmanuel, Noel Crespi and Thomas Magedanz, eds. 2013. "Evolution of Telecommunication Services: The Convergence of Telecom and Internet: Technologies and Ecosystems." Springer-Verl Lecture Notes in Computer Science. Heidelberg, Germany: Springer.
- Blind, Knut. 2004. *The Economics of Standards: Theory, Evidence, Policy*. Cheltenham, UK: Edward Elgar Publishing.
- Bohlin, Par Erik and Stanford L. Levin, eds. 1998. *Telecommunications Transformation: Technology, Strategy and Policy*. Amsterdam, The Netherlands: IOS Press.
- Chen, M., S. Mao, Y. Zhang and V. C. Leung. 2014. *Big Data — Related Technologies, Challenges and Future Prospects; Chapter 5, "Big Data Analysis."* Springer Briefs in Computer Science. Heidelberg, Germany: Springer.
- Delic, Kemal A. 2016. "IoT Promises, Perils and Perspectives." *ACM Ubiquity*, February, 1–5. doi: 10.1145/2822889.
- GECON. 2007. "Grid Economics and Business Models: 4th International Workshop, GECON 2007 Proceedings." Rennes, France: Springer-Verlag Lecture Notes in Computer Science.
- ISOC. 2015a. "Open Internet Standards: An Internet Society Public Policy Briefing." October 30. www.internetsociety.org/sites/default/files/ISOC-PolicyBrief-OpenStandards-20151030.pdf.
- Xiaoguo Zhu, Kevin and Zach Zhizhong Zhou. 2011. "Lock-In Strategy in Software Competition: Open-Source Software vs. Proprietary Software." *Information Systems Research* 23(2): 536–545. doi: 10.1287/isre.1110.0358.

ABOUT THE AUTHOR

Patrik Fältström is the head of Engineering, Research and Development at Netnod. He has previously served in leadership roles at IETF, the Internet Architecture Board and Internet Society (ISOC) (1998–2009); he was a distinguished engineer at Cisco Systems and has worked at Tele2 in Sweden, the Swedish Navy and Bunyip in Canada. Patrik was also appointed adviser to the Swedish IT minister (2003–2014), and has been the chair of the Internet Corporation for Assigned Names and Numbers Security and Stability Committee since 2011. Patrik has been working with Internet-related standardization and development since 1985 and is one of the founders of the ISOC Special Interest Group on Internet of Food.

Patrik holds an M.Sc. in mathematics from the University of Stockholm. He is a member of the Royal Swedish Academy of Engineering Sciences. In January 2011, Patrik received the Order of the Cross of Terra Mariana, V class, from the president of Estonia.

**CHAPTER NINE:
WHEN ARE TWO NETWORKS BETTER THAN ONE?
TOWARD A THEORY OF OPTIMAL FRAGMENTATION**

Christopher S. Yoo

Copyright © 2016 by the Centre for International Governance Innovation and the Royal Institute of International Affairs

INTRODUCTION

The Internet has made it possible for the world's citizens to connect with one another and access information to an unprecedented extent. The convergence of much of the world's communications media and information onto a single platform has yielded benefits that were unimaginable a few decades ago.

But a number of high-profile recent developments have raised concerns among the Internet community that the Internet could fragment. Countries such as China have long asserted control over the content that their citizens can reach. Edward Snowden's revelations about the level of surveillance being conducted by the US government has led to calls for laws requiring that all data associated with a country's citizens be hosted domestically and that companies use only domestically produced software. Other commentators have criticized commercial practices that create fragmentation by favouring some types of Internet traffic over others.

If implemented, such policies could cause the benefits long associated with the Internet to attenuate or dissipate. Standardization, on the other hand, creates real benefits in ensuring that both consumers and producers can reach one another through a common platform, regardless of location, technology or application.

Acknowledging the benefits associated with widespread connectivity and interoperability does not necessarily lead to the conclusion that standardization and interconnection are always preferred over any form of fragmentation. Indeed, some networks running the Internet Protocol (IP) are not interconnected with the rest of the network, and others operate on somewhat different principles. Unless all decisions that are not standardized are simply assumed to be mistakes, the persistence of these networks suggests the existence of considerations cutting in the other direction that need to be understood before all fragmentation is categorically prohibited.

This chapter's central claim is that current discourse suffers from two basic problems. The first is that fragmentation is a reality; indeed, as described in the opening section, below, each major area of the Internet is already fragmented. The second is that the discourse has not offered a basis for determining whether and when fragmentation is good or bad. If fragmentation is always detrimental, as some seem to suggest, the optimal outcome would be for every network in the world to interconnect and to operate on a single, unified standard. The fact that this is not the case invites some exploration of the forces tending to favour unification and the forces tending to favour fragmentation as a basis for determining optimal network size.

CURRENT EXAMPLES OF FRAGMENTATION

Four different types of fragmentation exist on the Internet: fragmentation of physical networks, of the address space, of protocols and of legal regimes.¹ The following section will describe each in turn and provide real-world examples of how each operates.

PHYSICAL NETWORKS

Commentators often stress the importance of having a single network through which everyone can reach everyone else. In fact, the need for a single unified telephone system was the primary rationale for ending the competitive era in US telephone service in the early twentieth century (Gabel 1969). The US Federal Communications Commission's (FCCs) 2010 Open Internet Order echoed this sentiment when it declared, "There is one Internet...that...should remain interconnected" (FCC 2010, 17934). The implication is that governments should mandate that all networks interconnect with one another on equal terms.

Despite this lofty rhetoric, a review of actual practices reveals that a large number of IP-based networks do not interconnect with the public Internet. A better understanding of the rationales underlying these practices reveals considerations against mandating universal connectivity.

Air Gaps

Perhaps the classic reason for a system not to interconnect with the public Internet is security. One of the standard practices for protecting system security is to block interconnection with other networks by maintaining an "air gap" between the system and the rest of the Internet. Such solutions are imperfect, as they can be bridged. For example, the Stuxnet virus that damaged Iranian centrifuges in 2010 may have been transmitted by an infected memory stick used by a Siemens employee to update software. Even though the fact that these networks are not interconnected to the Internet as a whole did not render these networks completely secure, many networks choose not to interconnect to the public Internet in order to reduce the likelihood of a security breach.

Private Networks

Private networks represent a more common example of non-interconnected networks. Although some are disconnected to maintain security, others remain private to connect high-volume access points in the most cost-effective manner. Still others isolate themselves from

¹ For a more recent paper providing a related taxonomy, see Drake, Cerf and Kleinwächter (2016). For another interesting exposition on fragmentation, see Huston (2015).

the rest of the Internet to guarantee quality of service by avoiding bandwidth sharing. A prime example is the financial services industry, which depends on microsecond latencies. Because the public Internet cannot deliver such speeds on a consistent basis, large parts of the financial services industry rely almost entirely on private networks.

Specialized Services for Voice and Video

Quality-of-service concerns also lead different providers to rely on segregated bandwidth to ensure delivery data associated with latency-intolerant applications, such as voice and video, in a timely manner. These providers sometimes dedicate capacity to these applications and do not make that capacity available for other users even when not used for voice or video. Although the channelization is often virtual, the bandwidth remains dedicated for single purposes and is not made available to other Internet users.

ADDRESS SPACE

Maintaining the unity of the address space is a long-standing principle of the Internet, dating back to its earliest days. Indeed, when Vint Cerf and Bob Kahn (1974) first articulated what would become the Internet suite of protocols, one of the central problems motivating their endeavour was the fact that different packet networks had different ways of addressing destination hosts.

Their solution was to create a uniform addressing scheme that could be understood by every network. Telephone networks once faced the same problem, and a unified numbering scheme has long been regarded as essential to maintaining universal reachability.

IPv6

Despite the widespread recognition of the benefits of unified address spaces, important counter-examples exist today. The transition to IP version 6 (IPv6) represents a prime example. As most observers are aware, the original IP version 4 (IPv4) header only allocated 32 bits in the header to the address space. That means that IPv4 can support just under 4.3 billion addresses.

At the time the Internet was created, 4.3 billion addresses seemed like more than the Internet would ever need. But the Internet has succeeded far beyond anything its creators ever imagined. As of November 2015, 3.4 billion of the world's 7.3 billion citizens accessed the Internet, and the rapid diffusion of data-enabled mobile phones and tablets means that individuals increasingly have more than one IP address. In addition, numerous businesses have Internet addresses as well. Moreover, industry observers predict that the advent of the Internet of Things will cause the connection of as many as 25 billion devices to the Internet by 2020.

The net result is that the Internet has run out of IPv4 addresses. To address this problem, the Internet Engineering Task Force created IPv6. The IPv6 header contains an address space that consists of 128 bits, which is enough to assign a separate address to every molecule in the solar system.

In making the transition from IPv4 to IPv6, network architects chose to make the IPv4 and IPv6 address architectures independent of each other, requiring each router to run both in parallel. This so-called dual-stack approach requires all IPv6 routers to implement parallel address structures simultaneously.

Middleboxes

Another deviation from the universal address space in which each machine has a unique address visible to all other users is the advent of middleboxes, such as network address translation (NAT) boxes. These devices temporarily mitigated the exhaustion of IPv4 addresses by allowing multiple hosts to share the same IP address by acting as if they were processes operating on a single machine instead of being distinct hosts. These devices are quite common; indeed, everyone who owns a Wi-Fi router uses them. They have also proven quite controversial in the technical community, because they deviate from the core principle of universal visibility of addresses and make it difficult, if not impossible, to reach certain parts of the network unless the person attempting to contact them has access to specialized state information or employs a NAT-traversal technique. The addition of middleboxes makes network operation more complicated and introduces per-flow state into the core of the network in ways that can make it less robust.

Proprietary Numbering Systems

A number of proprietary numbering systems have emerged in voice over IP. The most important of these is Skype, which provides unique addresses to allow users to connect with other Skype users for free. Skype also interconnects with the public telephone system, so that Skype users may also use traditional telephone numbers for a fee to contact non-Skype users. But the fact remains that Skype users have two parallel, non-interconnected address structures.

The Domain Name System

Address fragmentation is also often raised with respect to the domain name system (DNS). Under the current architecture, the assignment of Universal Resource Locators and IP addresses is done on a distributed basis. Coordination of these different hierarchies of names and numbers depends on the fact that they all refer back to a common root file that determines the authoritative name servers for each top-level domain. The historical role the US government has played in creating the Internet left the

Commerce Department with veto power over any changes to the root zone. The Edward Snowden revelations raised serious concerns about the role of the US in Internet governance, which in turn has led some countries to consider shifting their reliance to different root files over which the US government has no control. The controversy over ongoing US oversight over the DNS led the Commerce Department to announce that it would transition oversight of these functions to a non-governmental entity. The current deadline for this transition is September 2016.

PROTOCOLS

When trying to connect heterogeneous networks, Cerf and Kahn faced more than just inconsistent address structures. The networks they were attempting to interconnect ran different protocols. They considered translating protocols every time packets crossed from one network to another. The problem was that translation introduces errors. Moreover, translation would have a difficult time operating at scale, as the addition of every new network protocol would require the reconfiguration of every other system attached to the Internet. Instead, as mentioned, Cerf and Kahn required that all networks connected to the Internet operate a single protocol, IP. Insisting that every system would recognize IP would guarantee universal connectivity.

During the network neutrality debate, many advocates have criticized the use of protocols to prioritize certain traffic over others. As an initial matter, what is commonly overlooked is that the Internet was designed from the beginning to support the ability to differentiate among different types of traffic. The need for routing policies can trace its origins to the acceptable-use restrictions prohibiting commercial traffic from traversing the original National Science Foundation Network. A review of the IPv4 header reveals that the designers included a type of service field intended to mark packets for particular kinds of prioritization. Subsequent changes have made this field more customizable. It remained sufficiently important to be retained during the transition to IPv6.

In addition, the most recent version of Border Gateway Protocol (BGP), which provides the basic routing functionality of the Internet, was designed to create routing policies. In other words, BGP was specifically engineered to allow different types of traffic to be treated differently. Although advocates of policies such as network neutrality argue that prioritization should never be used, despite the fact that such functionality has been part of the Internet's design from the outset, using it does not necessarily represent fragmentation of protocols.

The more fundamental problem is that no one protocol does everything well, and every protocol necessarily involves trade-offs. IP is no exception. Although it has proven incredibly robust, the engineering literature is replete with acknowledgements of functions that the current Internet

does not perform well. These include security, mobility, "multihoming", video distribution and cost allocation, to name a few.

While these shortcomings were not that important when the Internet was largely about email and web browsing, in the modern Internet these new functions are now mission critical. This is causing pressure to evolve new protocols. In fact, the Internet is operating a number of protocols that are not completely consistent with the Internet approach. One example is Multiprotocol Label Switching (MPLS), which routes on the basis of specialized labels instead of IP addresses and employs an approach that bears some aspects of circuit switching. Such protocols are widely used to provide the functionality for voice and video that the traditional Internet cannot support. Only firms sharing access to the flow labels associated with MPLS can route traffic associated with these flows. Moreover, firms are using MPLS to implement a wide range of routing policies.

Interestingly, although MPLS initially represented fragmentation of the protocol space, recent changes appear to have incorporated it back into IP architecture. During the transition to IPv6, designers greatly simplified the architecture by removing a large number of fields from the IP header. The one field they added was to introduce a flow label field. This change effectively makes MPLS consistent with the basic architecture rather than representing an example of fragmentation.

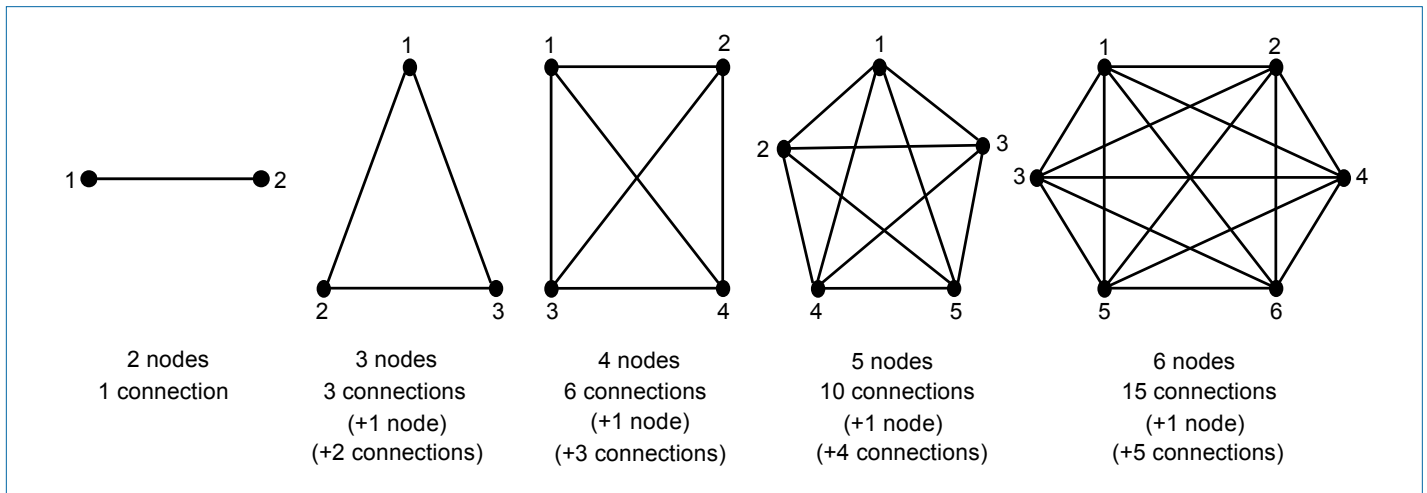
The fact that the Internet architecture has now evolved to incorporate MPLS should not overshadow the larger story. Network features emerge that fragment basic architecture. Some will be incorporated into the design, but they represent deviations during the interim. Still other innovations will never be assimilated into the design.

LEGAL REGIMES

The law has also struggled with the question of optimal fragmentation. Clearly, having the same conduct treated differently by multiple jurisdictions can impose a burden on commerce. Indeed, the desire to harmonize law and to reduce internal barriers to trade was one of the key purposes underlying the founding of the United States, and remains one of the central goals of the European Union.

At the same time, many important areas of the law in America have been left to state jurisdiction. For example, privacy law on data breaches is largely a matter of state law, and reporting requirements remain quite different across the United States and in the European Union. Contract law governing e-commerce is a creature of state law. Taxation remains entirely within the control of states, as does criminal law. Furthermore, in the European Union, individual rights remain a matter of member states' national law. In short, many areas of law are subject to considerable fragmentation.

Figure 1: The Relationship between the Number of Nodes and the Number of Connections



Source: Author.

THE BENEFITS OF UNIFICATION


The primary argument against fragmentation is based in the economics of network effects. Network effects exist when the primary determinant of a network's value is the number of other users connected to the network. The more people that an individual subscriber can reach through the network, the more valuable the network becomes, even when the nature of the service and the price paid for it remain the same.

The theoretical basis for network economic effects is known as Metcalfe's law, named after Robert Metcalfe, the inventor of the Ethernet, who first highlighted the importance of this relationship.² Metcalfe's law is based on the insight that as a network grows in size, the number of potential connections increases faster than the number of nodes.

Stated more generally, if the number of nodes equals n , the number of potential connections equals $(n^2 - n)/2$, which means that the number of potential connections increases quadratically with the number of nodes. This causes the number of connections to increase very rapidly. For example, the first 100 nodes create almost 5,000 potential connections. Adding another 100 nodes (to 200) increases the number of potential connections to just under 20,000, an increase of nearly 15,000. Adding yet another 100 nodes (from 200 to 300) increases the number of potential connections to almost 45,000, an increase of nearly 30,000. Further additions by increments of 100 nodes will cause even larger increases in the number of potential connections.

² For discussions on the connection between Metcalfe's law and network economic effects, see Yoo (2012; 2015).

Figure 2: The Systematic Value of Compatibly Communicating Devices Grows as the Square of Their Number



Source: Metcalfe (2006).

Assuming that each potential connection increases the value of the network by an equal amount, increases in network size cause a quadratic increase in network value. Assuming that the cost of adding nodes is constant, increases in network size cause a linear increase in cost. The result is inexhaustible returns to scale, in which bigger is always better, as demonstrated by the figure Metcalfe used to communicate the concept during the early 1980s (reproduced above). Metcalfe's law is widely celebrated as the foundation of the Internet's success.

OFFSETTING CONSIDERATIONS

Metcalfe's law provides a clear theoretical basis for opposing the fragmentation of networks. The larger a network becomes, the greater the surplus between benefits and costs. Indeed, if Metcalfe's law were the only consideration, every part of the industry would consist of a single network. But as shown in the opening section, many parts of the industry consist of multiple fragmented networks. This makes it important for us to identify those

factors pushing against the tendency for networks to combine into a single large network.

DIMINISHING MARGINAL RETURNS

Consider first the assumption that increases in the number of potential connections cause quadratic increases in network value. This presumes that potential connections increase value no matter how many links have already been established.

Empirically, such a result is quite unlikely. As a seminal article on network economic effects has noted, because those who place the highest value on the network are most likely to be the first to adopt, one would expect later users to provide less value (Rohlf's 2001, 29). A moment's reflection undercuts the expectation that adding more connections would continue to add value. If an end user has access to only one auction or sports news site, the marginal value of adding another is very high. If, on the other hand, the end user already has access to one hundred different versions of each type of site, the incremental value of having access to another version is much smaller.

Bob Briscoe, Andrew Odlyzko and Benjamin Tilly (2006) have provided the most sophisticated critique of Metcalfe's law. They argued that the inexhaustibility of the returns to scale is the direct result of the assumption of the model. They argue that the diminishing marginal returns inherent in the network are better captured by a rule of thumb known as Zipf's law, which holds that if some large collection of elements is ordered by size or popularity, the second element in the collection will be about half the measure of the first one, the third one will be about one-third the measure of the first one, and so forth. Stated more generally, the value of the n th item in the collection will be $1/n$ of the first item. In other words, the value of additional items decays exponentially.

Metcalfe (2006) responded by arguing that Briscoe, Odlyzko and Tilly had misunderstood his work. His point was not to assert that returns to scale in network size are inexhaustible, but rather to underscore how the adoption of a network depended on reaching a critical mass of users. In a later publication, Metcalfe (2013) recognized that the effect he described would not apply to very large networks: "Metcalfe's Law might overestimate the value of a network for a very large N . A user equipped to communicate with 50 million other users might not have all that much to talk about with each of them. So maybe the growth of systemic network value rolls off after some N ." Metcalfe pointed out that inexhaustible returns to scale were also a feature of the Zipf's law approach advocated by Briscoe, Odlyzko and Tilly. He also presented empirical evidence based on Facebook usage, suggesting that Metcalfe's law represented a better measure of value than Zipf's law.

Ultimately, the impact of diminishing marginal returns is an empirical question. Fortunately, the claim advanced here does not depend on resolving who has the better of the argument. It suffices to point out that circumstances may exist where further increases in network size will not yield substantial value. This means that whether fragmentation or unification is the better strategy is a question that must be studied empirically, not merely be asserted.

THE VALUE OF HETEROGENEITY

Complementary to the problem of diminishing marginal returns is the fact that people may place different absolute value on different potential connections. In other words, the locations that end users frequent on the Internet are not randomly distributed across the entire Internet. Instead, end users typically focus their visits on a small number of locations.

For example, empirical studies have shown that the average telephone user exchanges calls more than once a month with only five other people (Galbi 2009). Studies of Facebook reveal that users similarly exchange personal messages with no more than four people per week and six people per month (Adams 2012). Indeed, Facebook patterns confirm a concept known as Dunbar's number, which suggests that the human brain can maintain no more than 150 close relationships at any one time (Dunbar 1993).

The result is that end users may not value the number of potential connections in the abstract as much as they value particular connections to specific locations. Speaking personally, my Internet usage is disproportionately concentrated on a handful of locations, including my office computer via remote desktop access, my email server, my bank and a handful of other financial institutions, a number of utilities for bill payment, and a few news sites and blogs. I would place a higher value on connectivity to the sites I visit the most than I would on the ability to connect to other locations.

Heterogeneity can also place pressure toward fragmentation in the context of protocols. The point is demonstrated eloquently in a simple paper authored by Joseph Farrell and Garth Saloner (1986), who wrote some of the pioneering papers on network economic effects. Assume that two different populations of end users (A and B) would each prefer a slightly different standard and that both would benefit from network economic effects if they were part of the same network. Each group has two options: It can join the other group's standard, in which case it gains from being part of a larger network, but loses value from adopting a standard that it prefers less. Or it can adhere to its preferred standard, in which case it benefits from consuming its preferred standard, but foregoes the benefits of network economic effects should the other group adhere to its preferred standard as well.

The considerations driving the equilibrium are clear. If the value that either group derives from consuming its preferred standard is sufficiently large, it will always adopt its preferred standard even if it means being part of a smaller network. Any welfare losses from network fragmentation are more than offset by gains in allowing groups of end users to consume a standard that is a better fit with their preferences.

Heterogeneity also explains fragmentation in law, as exemplified by the debate over federalism. In any federal system, an issue may be addressed at the federal level or at the regional level. Two of the primary reasons to address issues at the regional level are a diversity of values and/or conditions, and the facilitation of experimentation. Allowing each regional jurisdiction to tailor its own solution allows the law to effect a better fit with local circumstances, but at a cost of greater legal fragmentation. Indeed, in the United States, many key areas of the law remain governed by state law and thus face fragmentation, including contract law, corporate law and data-breach notification. Admittedly, there are other considerations tending toward federal resolution, including externalities, scale economies, “races to the bottom,” expertise and the potential of interest-group capture. Whether fragmentation or unification is the optimal outcome depends on which of these considerations dominates.

DETERMINING VALUE BASED ON THE TOTAL NUMBER OF NODES INSTEAD OF THE TOTAL NUMBER OF UNIQUE CONNECTIONS

Metcalf’s law also presumes that the value of a network is determined by the total number of unique pairwise connections. In other words, the value is determined by the ability to reach specific people. One can easily imagine situations where value depends on the total number of people one can reach through the network without placing any value on the ability to reach discrete people. A prime example of this is advertising. Many advertisers do not care if they can reach any particular person. Instead, they care only about the total size of audience.

Shifting the focus to the total number of people a network can reach without placing any value on the ability to reach particular individuals fundamentally changes the underlying economics (Nuechterlein and Yoo 2015). The fact that advertising represents the dominant source of revenue on the Internet suggests that heuristics such as Metcalf’s law may well overstate the value of preventing any action that may cause the network to fragment.

NONLINEAR INCREASE IN COSTS

Metcalf’s law also depends on the assumption that costs would increase linearly, which in turn is based on the assumption that the equipment costs of adding each

additional node would be precisely the same. The problem with this assumption is that there are other important sources of costs in the Internet.

The most important source of costs is congestion. The Internet is a shared medium. Indeed, the ability to multiplex streams of data across the same connection is one of the primary advantages associated with packet switching. Like any shared medium, the Internet can become congested if too many people attempt to use it at the same time. As congestion becomes severe, the costs grow much faster than linearly. Indeed, when buffers become completely full, the network can suffer from complete and sharply discontinuous lockout.

Another problem associated with the growth of the Internet is search costs. As more nodes are attached to the network, those who wish to use the network must incur higher search costs to find content that fits their preferences. The problems associated with this have led some to question whether certain social networks, such as Facebook, have become too big.

REAL-WORLD SOLUTIONS TO FRAGMENTATION

The presence of opposing considerations provides a framework for evaluating when unification is the optimal approach and when fragmentation might yield benefits. As such, it also provides a basis for describing the world as it exists today, in which some matters are unified or addressed at the federal level and others are fragmented or handled at a regional level.

Even when fragmentation exists, both engineering and law have developed institutions to manage the heterogeneity. The most important of these are partial compatibility and informal harmonization.

PARTIAL COMPATIBILITY THROUGH GATEWAYS

One way that networks can mitigate the problems associated with fragmentation is through gateways (also called adapters or converters) between networks. Many of the leading scholars on network economic effects have shown that perfect gateways can completely mitigate the problems of fragmentation (Matutes and Regibeau 1988; David and Bunn 1988; Katz and Shapiro 1994; Farrell and Saloner 1992). Farrell and Saloner further showed that even if the gateway is imperfect, it can mitigate the problems of incompatibility in whole or in part. Such gateways can ameliorate potential fragmentation in the physical architecture, the address space and the protocols.

ARBITRATION

For legal fragmentation, the most prominent means of harmonization is the resort to commercial arbitration. Commercial arbitration is honored internationally now by almost every jurisdiction and allows parties to opt in to a unified legal regime. Indeed, an arbitration clause can avoid national jurisdiction by opting to be bound by a pre-existing body of arbitral precedents.

CONCLUSION

Debates about Internet fragmentation often take on an alarmist tone that intimates that any practice that introduces a degree of heterogeneity into the network must be stopped. If followed to its logical conclusion, this point of view would mandate that all networks interconnect with one another on equal terms and operate the exact same protocols to ensure maximum interoperability.

The pragmatic perspective that animates network engineering generally regards such absolutist perspectives with suspicion. Often, multiple forces push particular outcomes in opposite directions. The natural response is to understand those forces and to study them empirically to determine how they should best be optimized. Undertaking such an analysis does not deny the value of wide-scale interoperability. There is no doubt that the “open Internet” standards have created tremendous benefits to the world and have proven more robust than anyone could have imagined.

The goals of this chapter are far more limited. It raises a defensive argument designed to raise the possibility that universal connectivity and interoperability may not be the preferred solution in every circumstance, and to try to identify heuristics to help guide the determination of when fragmentation is bad and when it might be good. Part of the argument is empirical: fragmentation and non-standardization are pervasive phenomena that exist in the physical network, the address space, the protocol space, and the law governing the Internet. Any evaluation of whether and when fragmentation is good or bad must seek to understand the forces that tend to push toward unification and toward fragmentation to help inform the proper balance in any particular case.

Finally, any assessment of fragmentation must take into account that participation in the Internet architecture is always voluntary. Those who operate IP-based networks always remain free not to interconnect them with the public Internet, to use different address structures or to use different protocols. Because interconnection and standards adoption remains voluntary, individual actors can be expected to interconnect or adopt the standard only when the individual benefits exceed the individual costs. Importantly, individual optimization decisions do not always lead to equilibria that are optimal for the network

as a whole. Thus, any assessment of fragmentation requires not only an understanding of when fragmentation and unification would be optimal globally, it also requires careful attention to the incentives of individual actors to determine whether the decentralized decision making that characterizes the Internet is likely to lead to good outcomes.

A version of this text was presented at the October 2014 meeting of the Global Commission on Internet Governance held in Seoul, Korea.

WORKS CITED

- Adams, Paul. 2012. *Grouped: How Small Groups of Friends Are the Key to Influence on the Social Web*. Berkeley, CA: New Riders Publishing.
- Briscoe, Bob, Andrew Odlyzko and Benjamin Tilly. 2006. "Metcalfe's Law Is Wrong." *IEEE Spectrum*. <http://spectrum.ieee.org/computing/networks/metcalfes-law-is-wrong>.
- Cerf, Vinton G. and Robert E. Kahn. 1974. "A Protocol for Packet Network Interconnection." *IEEE Transactions on Communications* 22 (5): 637–48.
- David, Paul A. and Julie Ann Bunn. 1988. "The Economics of Gateway Technologies and Network Evolution: Lessons from Electricity Supply Industry." *Information Economics and Policy* 3 (2): 165–202.
- Drake, William, Vinton Cerf and Wolfgang Kleinwächter. 2016. "Internet Fragmentation: An Overview." World Economic Forum Future of the Internet White Paper. www3.weforum.org/docs/WEF_FII_Internet_Fragmentation_An_Overview_2016.pdf.
- Dunbar, R. I. M. 1993. "Coevolution of Neocortex Size, Group Size and Language in Humans." *Behavioral and Brain Sciences* 16 (4): 681–735.
- Farrell, Joseph and Garth Saloner. 1986. "Standardization and Variety." *Economics Letters* 20 (1): 71–74.
- . 1992. "Converters, Compatibility, and the Control of Interfaces." *Journal of Industrial Economics* 40: 9–35.
- FCC. 2010. "Preserving the Open Internet," Report and Order. *Federal Communications Commission Record* 25:17905–18098.
- Gabel, Richard. 1969. "The Early Competitive Era in Telephone Communications, 1893–1920." *Law & Contemporary Problems* 34 (2): 340–59.
- Galbi, Douglas. 2009. "Telephone Social Networks." *Purple Motes*. <http://purplemotes.net/2009/11/29/telephone-social-networks>.
- Huston, Geoff. 2015. "Thoughts on the Open Internet – Part 2: The Where and How of 'Internet Fragmentation.'" CircleID. www.circleid.com/posts/20151006_open_internet_part_2_where_and_how_of_internet_fragmentation/.
- Katz, Michael L. and Carl Shapiro. 1994. "Systems Competition and Network Effects." *Journal of Economic Perspectives* 8 (2): 93–115, Spring.
- Matutes, Carmen and Pierre Regibeau. 1988. "'Mix and Match': Product Compatibility Without Network Externalities." *RAND Journal of Economics* 19 (2): 221–34.
- Metcalfe, Bob. 1995. "Metcalfe's Law: A Network Becomes More Valuable as It Reaches More Users." *InfoWorld*, October 2.
- . 2006. "Metcalfe's Law Recurses Down the Long Tail of Social Networking." VCMike's Blog. <https://vc mike.wordpress.com/2006/08/18/metcalfe-social-networks/>.
- . 2013. "Metcalfe's Law After 40 years of Ethernet." *Computer*, December.
- Nuechterlein, Jonathan and Christopher S. Yoo. 2015. "A Market-Oriented Analysis of the 'Terminating Access Monopoly' Concept." *Colorado Technology Law Journal* 14: 21–36.
- Rohlf, Jeffrey H. 2001. *Bandwagon Effects in High Technology Industries*. Cambridge, MA: MIT Press.
- Yoo, Christopher S. 2012. "When Antitrust Met Facebook." *George Mason Law Review* 19 (5): 1147–62.
- . 2015. "Moore's Law, Metcalfe's Law, and Optimal Interoperability." *Colorado Technology Law Journal* 14 (1): 87–102.

ABOUT THE AUTHOR

Christopher S. Yoo is a senior fellow with CIGI's International Law Research Program. He is based at the University of Pennsylvania, where he is the John H. Chestnut Professor of Law, Communication, and Computer & Information Science, and the founding director of the Center for Technology, Innovation & Competition. He was previously a professor of law at Vanderbilt University in Tennessee and has held visiting academic appointments in Australia, China, Germany, Italy, Japan, South Korea and Switzerland. Before entering the academy, he clerked for Justice Anthony M. Kennedy of the Supreme Court of the United States.

CHAPTER TEN: A FRAMEWORK FOR UNDERSTANDING INTERNET OPENNESS

Jeremy West

Copyright © 2016 by the Organisation for Economic Co-operation and Development

ACRONYMS

CGN	carrier-grade NAT
DNS	domain name system
http	hypertext transfer protocol
IP	Internet Protocol
NAT	network address translator
OECD	Organisation for Economic Co-operation and Development
RTBF	right to be forgotten
SMTP	Simple Mail Transfer Protocol
TCP/IP	Transmission Control Protocol/Internet Protocol
UDP	User Datagram Protocol
URL	universal resource locator

INTRODUCTION

“As divergent forces tug at the internet, it is in danger of losing its universality and splintering into separate digital domains,” *The Economist* (2010) stated. That was now more than five years ago. Although the Organisation for Economic Co-operation and Development ([OECD] 2008; 2014), among other bodies, has recognized the link between a distributed, interconnected architecture designed to be open by default and the Internet’s catalyst role for economic growth and social well-being, the splintering forces remain. These forces vary widely in nature and apply pressure at different levels of the Internet. They can be found in private sector actions as well as in public policies and governance.

A key question for policy makers is where they should aim to position their countries in the multidimensional Internet openness space. A number of important multi-stakeholder objectives — for example, sovereignty, public safety and economic development — call for actions that can lead to different degrees of openness. Because the Internet is a “network of networks,” the probability that interventions will have unintended consequences is higher than it would otherwise be. Addressing the needs of some stakeholders could be politically expedient, for example, but it might also cause unintended harm to the more numerous but less visible masses. Setting and implementing sound policies related to openness can therefore be a challenging undertaking.

To help policy makers reach more informed decisions about Internet openness, the OECD has begun to develop a framework for analysis. It includes a definition of Internet openness, a broad description of the types of benefits — as well as some of the harms — that are associated with it, and a suite of relevant stakeholder objectives. The OECD is also looking at the way those objectives are

translating into actions and conditions, with particular attention to how they affect openness at different layers of the Internet. The scope of the OECD’s project includes gathering initial evidence of the economic and social benefits of Internet openness (and the impact of reducing openness), with a focus on international trade, innovation and entrepreneurship, macroeconomic performance and societal well-being. This chapter, drawing on research conducted by the author for the OECD Committee on Digital Economy Policy, proposes a definition of “Internet openness.”¹

THE OPEN INTERNET VERSUS INTERNET OPENNESS

Although the term “open Internet” is used frequently, it has no universally accepted definition. It is a convenient phrase, like “level playing field,” that glosses over complexities. It tends to be used on the assumption that everyone agrees on its meaning, but they do not. To some, it means technical openness (for example, global interoperability of transfer protocols). To others, it means openness in a human rights sense (such as freedom from online censorship). Many use it interchangeably with other terms that do not have a universally adopted definition (for example, “net neutrality”), or it may be intended as shorthand for a particular characteristic such as geographically or demographically broad access to the Internet. As a result, the term causes confusion.

Furthermore, speaking about an open Internet suggests that Internet openness is binary — that it can only be fully open or fully closed. Even if one considers only the technical aspects of openness, the binary view does not correspond with how the Internet actually works. The Internet is a layered arrangement consisting of a physical access and transport infrastructure, an agreed set of packet and transport protocols, a domain name system (DNS), an Internet Protocol (IP) address system, applications and content. Together, the layers enable data flows that travel between user devices located at the edges of the network. Technical openness depends on the conditions at each of those layers. Some of the conditions increase openness, while others restrict it. Some even do both simultaneously. Certain conditions affect openness more strongly than others, and they can also affect different aspects of openness. But they do not simply turn openness “on” or “off.”

For example, one condition that affects openness at the IP address layer is the shortage of IP addresses that has arisen due to the limitations of IP version four (IPv4), a protocol that identifies devices on a network. The shortage of available IP addresses makes it harder to connect

¹ This chapter should be read in conjunction with the chapter *Internet Openness and Fragmentation: Toward Measuring Economic Effectiveness*, by Sarah Box.

more users and devices to the Internet (a closing effect). Therefore, a workaround solution — called a network address translator, or NAT — was created. A NAT allows multiple devices to share the same IP address. Many of the boxes that provide fixed broadband Internet access and Wi-Fi in homes have NATs built into them, enabling all of the Internet-connected devices within the home to use the same IP address. A carrier-grade NAT, or CGN, is a supersized NAT that allows many homes and other end sites to share small pools of IP addresses. CGNs increase openness by improving access to the Internet. However, they do not provide unlimited access, and in any event CGNs simultaneously reduce accountability by essentially hiding or anonymizing user activity — a closing effect. Consequently, CGNs neither fully open nor fully close the Internet, but they do affect its openness.

In fact, the Internet has rarely, if ever, been either fully open or fully closed. On the one hand, absolute openness — if such a state is even possible — would require the end of arrangements that are critical for economic and social reasons, such as having to pay for hardware and Internet access and enforcing child pornography laws. On the other hand, total closure would transform the Internet into nothing more than a series of isolated nodes, at which point it would cease to be a network at all.

The reality is that the Internet has degrees of both openness and “closedness” along many vectors. Therefore, the question to ask is not whether the Internet is open or closed, but how much openness or closedness it has, and in what dimensions. In fact, Internet openness is always in a state of flux, continuously becoming more open in some dimensions and more closed in others.

Accordingly, it is more helpful to study Internet openness with a multidimensional space in mind than with a basic open-or-closed perspective. That is why the oversimplifying term “open Internet” has been rejected in this chapter in favour of “Internet openness.”

In keeping with that choice, this chapter adopts a broad view of Internet openness, one that goes well beyond a purely technical view and encompasses economic, social and other factors. On the one hand, technical openness increases when openly available protocols are used consistently to receive and send data flows across interoperable layers of the Internet, relying on an open and consistent IP address system and a uniform convention for domain names. Thus, for example, the more that devices connected to the Internet consistently use the Transmission Control Protocol/Internet Protocol (TCP/IP), the more technical openness there will be. On the other hand, the more that non-standard data flow control algorithms are used, the less technical openness there will be.

Economic openness varies with the ability of users to get online and to use the Internet to enhance their economic

opportunities and to put them to productive uses. For instance, economic openness increases as broadband infrastructure grows, but it decreases when access providers lack competition and charge higher prices or provide poorer service as a result.

Social openness is positively related to the ability of individuals to use the Internet to broaden their non-pecuniary opportunities, such as keeping in touch more easily with family and friends, becoming more informed about topics of interest to them or expressing themselves. As an illustration, social openness increases when laws curtailing political expression are eased. It decreases when access to online educational material is eliminated because a government decides to block the entire platform through which the material is available.

OPENNESS AT A GLANCE

Table 1 sets out the elements of openness that are discussed throughout this chapter.

TECHNICAL OPENNESS

A core feature of technical openness is the end-to-end principle (Saltzer, Reed and Clark 1981; Blumenthal and Clark 2001). The intended role of an open switched network that follows the end-to-end principle is limited to carrying individual data packets from source to destination. It does not alter or interfere with the packets; it just transports them, and it does so without favouring one stream of packets over another. All user access and all functions and services that populate the network are provided by devices that sit outside of the network itself. These devices communicate among themselves in a manner that is largely opaque to the network. In other words, the network should not replicate functions that can be performed by communicating end systems.

Like most elements of openness, the end-to-end principle is not an all-or-nothing absolute requirement, though. Rather, it is a principle that, in practice, may be followed to a greater or lesser degree in a network. The more it is followed, the more openness the network has. Stakeholders may thus prefer, or aspire to, an ideal of a fully end-to-end network, but just because a network might not be 100 percent end-to-end in practice does not mean that there is no openness in the network. Thus, the end-to-end principle is not to be confused with a set of network engineering constraints. Various services may operate in ways that are not precisely aligned to it. However, the extent to which particular network components can successfully operate while not adhering exactly to these broad precepts is bounded by the ability of other network components that operate according to these principles to successfully interoperate with them.

Table 1: Elements of Internet Openness

Technical	Economic	Social	Other
<ul style="list-style-type: none"> • End-to-end principle: <ul style="list-style-type: none"> – use of consistent standards – interoperable – open, consistent address space – uniform convention for domain names • Open protocols for core functions 	<ul style="list-style-type: none"> • Cross-border supply and consumption • Economic accessibility • Regulatory transparency and certainty 	<ul style="list-style-type: none"> • Respect for human rights: <ul style="list-style-type: none"> – freedom of expression – freedom to associate – privacy – freedom from discrimination – education 	<ul style="list-style-type: none"> • Digital security: <ul style="list-style-type: none"> – availability – integrity – confidentiality – <i>but</i> with some vulnerability • Empowerment of users over data sent and received • Distributed control • Inclusive governance • Multilingualism

Source: Author.

In an open switched network, the end-to-end principle requires the use of consistent technical standards. That means all active, packet-switching elements in the network use a uniform interpretation of the contents of each packet, supporting precisely the same protocol (in the case of the Internet, this is the IP specification). Consistency also means that all connected systems inside the network are able to communicate by using the same transport protocols. The Internet has commonly adopted two end-to-end transport protocols, the TCP and the User Datagram Protocol (UDP). While many other transport protocols have been defined, common convention in the Internet has settled on TCP and UDP as the two “universal” end-to-end transport protocols. The more consistently that connected systems around the world communicate by using these protocols, the more Internet openness increases.

Consistent technical standards contribute to another feature of technical openness: interoperability, that is, the ability to use any layer of the Internet without arbitrary, technical restriction. (Such use is not necessarily free of charge, however.) Furthermore, interoperability implies that there are no inherent or arbitrary technical restrictions interfering with anyone’s ability to provide goods and services at any layer, whether it be transmission capacity, switching, domain names, applications or any of the other layers that make up the Internet. Interoperability leads to greater freedom of choice: the freer consumers are to choose the devices, applications and services they use, and the freer providers are to choose the types of devices,

applications and services they offer, the more open the network is deemed to be.²

The end-to-end principle also demands an open, consistent address space. This condition means every destination on the Internet is reachable from any other location on the Internet, which requires all destinations to have their own IP address that everyone else can reach. IP addresses must therefore be allocated and administered in such a way that each address is uniquely associated not only with a single network, but with a single device within that network. The network itself cannot resolve clashes where two or more devices are using the same address, so the responsibility for ensuring that all addresses are used in a manner that is unique is left to the bodies who administer address allocation and registration.³

The next requirement of the end-to-end principle is a uniform convention for domain names. The DNS is the combination of a common convention for creating names and a consistent methodology for transforming a

2 Note that “interoperability,” as the term is used here, refers to interoperability with the network. It does not imply that devices sitting outside the network must be interoperable with each other, but only that the protocols used by the network should be available to device makers so that they can make their products compatible with the network. Thus, for example, iPhones and Android phones can both connect to the Internet, but they run on different operating systems.

3 Address allocation and registration has been an evolutionary process. The original address administration and registry function was managed through US research agencies. The evolution of that model led to the creation of five regional Internet registries, each of which serves the address allocation and registry function needs of regional communities. The practices relating to access of address space through allocation and assignment are based on policies developed by the respective address communities in each region. The general theme of these policies is one of “demonstrated need,” where addresses are available to applicants who can demonstrate their need for these addresses within their intended service infrastructure.

universal resource locator (URL) from a format that is easy for humans to use into a format that is easy for machines to use (the “name resolution” function). In other words, the DNS allows people to use familiar symbols and terms, such as “www.oecd.org,” when referring to service points connected to the Internet, instead of numeric IP addresses and transport protocol port numbers, such as “194.66.82.11.” For the DNS to work properly, certain rules have to be followed when creating the names, and each name has to be tied to a single IP address.

Whenever data is sent from one Internet-connected device to another, there is a DNS query. The query asks the DNS what the correct IP address is for the desired recipient of the data flow. Regardless of where and how a DNS query is generated, the response should reflect the current state of the authentic information published in the DNS. The implication here is that the DNS uses the name space derived from a single and unique root zone, with all name resolvers answering name queries by searching within that uniquely rooted name space. If that does not occur, then, when a user types, for example, “www.yahoo.fr” he or she might wind up looking at the home page for, say, *El País*, thereby introducing an element of chaos that would severely undermine the Internet’s utility.

The more closely and consistently the end-to-end principle is followed, the greater the likelihood that no matter where data originates and what path it takes as it travels across the Internet, it will arrive intact at the intended destination, and only that destination.

Finally, technical openness also increases with the adoption of open protocols, at least for a number of core Internet functions. Open protocols are openly available, meaning they are not encumbered by restrictive claims of control or ownership. A number of open, commonly defined application-level protocols have already been adopted for core services. For example, applications that pass email messages are expected to use the Simple Mail Transfer Protocol (SMTP) and browsers are expected to use the hypertext transfer protocol (http). Other network-wide functions, including data transfer, instant messaging and presence notification, are also supported by open protocols.

However, proprietary protocols do exist, even for core functions such as sending data across the Internet. Some companies have incentives for using proprietary transit protocols. Their motive, at least in some instances, is to try to use a disproportionate share of the available bandwidth for their own communications without experiencing packet loss (which occurs when packets of data travelling across the Internet do not reach their destination). See Box 1 for more details.

The open nature of the technical foundation of the Internet is critical to the Internet’s “identity.” It is what it is today

largely because of its technical openness. Policy actions and inactions that restrict technical openness have the capability to weaken the Internet’s security, flexibility and stability.

ECONOMIC OPENNESS

The Internet’s economic openness corresponds to the ability of people, businesses and organizations to get online and use the Internet to increase their economic opportunities and capitalize on them. Increasing one’s economic opportunities via the Internet naturally depends on access to the Internet. Having economic access means that the requisite infrastructure for connecting to the Internet is available, and at a competitive price. The better the markets for Internet service, computers, smartphones and other connecting devices function, the more open and inclusive the digital economy will be. Economic access requires investment in electricity and broadband infrastructure as well as sound competition policy (OECD 2014, 7, 19-20).

Consider the case of telecommunications market liberalization in Kenya. When Telkom Kenya’s monopoly on the Internet backbone ended and two new firms entered the scene, they brought competition into the country’s market for Internet access for the first time. As a result of that and other pro-competitive policies, bandwidth availability increased and service costs to operators declined. In fact, their rates dropped by some 90 percent and those savings were passed along to consumers, who also benefited from wider geographic access. The number of Internet users in Kenya more than doubled during the year after liberalization. “Today, thanks largely to a liberal market approach complemented by proactive and effective policymaking, Kenya is a regional hub for tech and Internet start-ups and has attracted substantial investment from employers like IBM and Microsoft” (Dalberg 2014, 18).

The access aspect of economic openness goes beyond merely being able to connect to the Internet. It also refers to the degree to which entrepreneurs — from individuals to global companies — can capitalize on the economic opportunities enabled by the Internet without interference from over-inclusive or anticompetitive regulations (for example, unnecessarily broad content-based filtering or blocking policies). Private sector conduct, such as making it unreasonably difficult to sell an application in a platform’s app store, can have a restrictive effect on economic openness, too. Conversely, the easier it is to legally use and sell applications, products, content and services on the Internet, the wider the economic opportunities will be.

Economic openness also refers to the ability to consume and supply services over the Internet on a cross-border basis. The fewer unjustifiable barriers there are that prevent users from accessing, generating and selling the lawful content, applications and services of their choice,

Box 1: Non-Standard Flow Control Algorithms

The end-to-end principle assumes that TCP is the predominant protocol used by hosts connected to the Internet. In particular, it assumes that the data flow control algorithm used by all TCP implementations behaves in very similar ways across the Internet. That algorithm relies on the aggregate outcome of the TCP flow control protocols to provide a fair-share allocation of common network resources, so that an approximately equal proportion of those resources is devoted to each active flow. In other words, no one flow is more important than any other.

Specifically, each TCP session will both impose pressure on and respond to pressure from other concurrent sessions in trying to reach a point where the network's bandwidth is shared equally across the concurrent active flows. Packet loss occurs when there is too much pressure, so a flow will gradually increase its sending rate until the onset of packet loss, at which point it will immediately halve its sending rate. It will then gradually probe with increased rates until the next packet loss event. TCP implementations that use a different flow control algorithm normally fare worse, as their efforts to put more pressure on other flows often result in packet loss in their own flow.

However, there has been a significant body of research into flow control algorithms and some have emerged that appear to be able to secure a greater relative share of network resources without the self-damage problem. These algorithms are capable of exerting "unfair" pressure on other concurrent TCP flows, consuming a disproportionate share of network resources. Examples include Akamai's FastDNS, Google's QUIC and some Linux distributions using CUBIC.

Source: Geoff Huston, consultant to the OECD.

and regulations concerning the Internet, and the fairer the process for enforcing them, the greater the regulatory transparency and certainty (OECD 2014, 10). Regulatory transparency and certainty increase economic openness by reducing one of the risks of doing business as either a buyer or a seller in the digital economy: the risk of violating applicable laws or of being unable to defend one's rights adequately.

SOCIAL OPENNESS

The Internet's social openness corresponds to the ability of individuals to use the Internet to broaden their non-pecuniary opportunities. Such opportunities could include their meeting new people and exchanging knowledge and ideas with them, keeping in touch more easily with family and friends, expressing themselves to a potentially wider audience than they would otherwise be able to reach, becoming more informed about topics that are personally meaningful, gaining a better understanding of what their elected representatives in government are doing and becoming more active in their communities. The social aspects of Internet openness can reverberate and have a positive effect on economic openness. In particular, enhancing elements such as freedom of expression promotes more than human rights; it promotes innovation, as well. Innovation depends greatly on knowledge sharing and collaboration, and restrictions on freedom of expression online can inhibit sharing and collaboration.

The protection, promotion and enjoyment of all human rights is closely connected to the Internet's social openness. Consecutive resolutions of the United Nations Human Rights Council affirm that all human rights apply online just as they do off-line. Human rights include, for example, freedom of opinion and expression, freedom to associate, privacy, and education (United Nations [UN] 1948, articles 12, 19, 20, 23, 26; UN 2012). To see how human rights can bear on social openness, consider freedom from discrimination (UN 1948, article 2), which is particularly relevant in the context of access. If individuals are being denied access to lawful content and services online on the basis of their race, colour, sex, language, religion, political or other opinion, national or social origin, and so on, there is an obvious negative effect on social openness. Conversely, then, the more access that individuals have to lawful content and services online without interference based on those factors, the more socially open the Internet is. (Interestingly, the relationship between human rights and Internet openness is mutually reinforcing. Not only does respect for human rights generally enhance openness, but openness facilitates human rights [OECD 2014, 20].)

Although the concept of Internet openness incorporates consideration of the respect accorded these rights, making human rights ever stronger will not necessarily always result in more openness. Eventually, some of these rights

regardless of the jurisdiction they are coming from or going to, the more economically open the Internet is considered to be (OECD 2014, 7). Examples of justifiable barriers to cross-border data (content) flows include well-tailored measures that protect public safety or preserve culture and national values. Note that privacy- and security-enhancing measures are not deemed to be barriers to openness when they balance fundamental rights, freedoms and principles and comply with the OECD's guidelines on privacy (OECD 2013) and security (OECD 2015). Indeed, such measures (discussed below) are considered to enhance openness.

Economic openness also depends on regulatory transparency and certainty. The clearer the laws, rights

would become so strong that they would impinge on each other and, as a result, on openness. For example, if freedom of expression were limitless, it would be legal to post child pornography on the Internet. See Box 2 for another example.

OTHER FACETS OF OPENNESS

Certain elements of openness do not fit neatly within the categories of technical, economic or social openness. They might cut across some or all of the categories, or they might just have different natures altogether. One such element is the empowerment of individuals to understand and control how their private data is used online, as well as to control the information they receive online (OECD 2014, 12). Empowerment corresponds with the degree to which Internet users are provided with useful, comprehensible information about the privacy ramifications of their online activities as well as the degree to which they can control those ramifications. Are there laws, regulations or industry codes of conduct in place that require online services to inform users about what personal data is being retained and how it will be used? To what extent do users have control over how their data is used? Note that in this context more openness for some stakeholders might imply less for others. For example, more openness for business in the form of greater freedom to use the personal data of its customers might imply less openness in the form of lower transparency, awareness or control for individuals. Conversely, more openness for individuals in the form of greater empowerment over their personal data might imply less openness for businesses.

The level of empowerment also depends on how much control users have over the amount and type of information they receive via the Internet. Are their email accounts flooded with spam? Are they able to block mail from certain accounts? Can they protect their children from content they consider to be harmful?

Empowerment is relevant to openness because it fosters trust in the Internet. The OECD's *Principles for Internet Policy Making* (2014, 25) envision a cooperative effort on empowerment, in which governments, the private sector, the Internet technical community and civil society "work together to provide the capacity for appropriate and effective individual control over the receipt of information and disclosure of personal data." The inclusion of the word "appropriate" reflects that a measured amount of control over one's personal data is called for.

Thus there can be too much or too little empowerment, but the right amount promotes openness. For example, great strides in medical research can be made with data that is collected via the Internet. If the data is suitably de-identified, the danger to personal privacy presented by its collection and use could be low while the benefits for human health could be high. However, if users were

Box 2: The Right to Be Forgotten

If extended far enough, some human or fundamental rights might eventually conflict with one another. For example, in 2014 the Court of Justice of the European Union ruled that under certain conditions individuals have the right to ask search engines to remove links with personal information about them. The right applies when the information is inaccurate, irrelevant, inadequate or excessive for the purposes of the data processing (Google Spain SL v Agencia Española de Protección de Datos, C-131/12, May 13, 2014, para. 93). The Court of Justice acknowledged that the "right to be forgotten" (RTBF) is not absolute and that it will therefore need to be balanced with other fundamental rights, such as freedom of expression (*ibid.*, para. 85).

The RTBF also illustrates the tension that can arise between privacy and openness. The RTBF increases privacy and therefore may increase trust, resulting in an opening effect. At the same time, the RTBF takes information off-line, which arguably has a closing effect. Each country must decide for itself how to manage the relationship between privacy and openness. Indeed, jurisdictions such as the European Union and the United States differ on the RTBF, as the right is protected in the European Union but not in the United States. Which jurisdiction has a more open Internet policy as a result is a subjective question.

able to invoke a blanket refusal that prevented any of their personal data from being used in any manner, no matter how many measures were taken to strip out its personally identifying tags, the result could well be considered a net loss for society.

Although Internet openness catalyzes a host of economic and social benefits, it can also expose users to online intrusions, fraud, extortion, ransomware, intellectual property theft, denial-of-service attacks and a variety of other dangers. Those cyber activities threaten economic and social well-being by exposing personal and private data, harming financial and public infrastructure, threatening public safety, subverting human rights and depriving businesses of the fruits of their innovation and investment. What is needed to combat these threats and to preserve the Internet's ability to carry global data flows safely is digital security. Security is, therefore, another element of openness. Security cuts across all of the dimensions — technical, economic and social — of openness, and has three main components.

CONFIDENTIALITY

The greater the availability to end-users of robust and uncompromised protection from third-party eavesdropping and unauthorized access to data, the more confidentiality they will have when they send and store data on the Internet (where “data” means any content that flows over the Internet, such as credit card numbers, bank account information, trade secrets, private conversations, photographs and so on).

INTEGRITY

The better able end-users are to verify the identity of whomever they are communicating with and to ensure that received communications are genuine and precise copies of what was sent, the more integrity their communications will have.

AVAILABILITY

The greater a network’s ability to withstand a cyber attack or hacking attempt without any interruption of service to users, the more availability that network has.

All else being equal, the more effective a network’s digital security measures are, the more users will trust and rely on the network. In short, any notion that digital security must be viewed as a closing element is incorrect, because it is critical for building trust in the Internet. If trust declines enough, people will be less likely to use the Internet than they would otherwise be and data flows will shrink. Consequently, a better way to look at digital security is to recognize it as an element that contributes to openness, provided it balances fundamental rights, freedoms and principles and complies with the OECD’s (2015) security guidelines.

This is not to say that absolutely airtight digital security would always be optimal (even if it existed, which it does not). Some degree of intrusion could be justified on grounds such as national security or law enforcement needs. In addition, stronger security comes at a financial cost, so it will be efficient for individuals and businesses to opt for a lower level of security for some or all of their activities.

Furthermore, any degree of Internet openness necessarily implies a certain amount of vulnerability. Internet security risks cannot be eradicated as long as the component networks remain interoperable and have any ability to communicate with one another. Ultimate security would require cutting oneself off from the Internet altogether, which would have an obvious closing effect. Accordingly, the OECD’s *Principles for Internet Policy Making* (2014, 11) recognize that “strong” privacy protection rather than “absolute” privacy protection “is critical to ensuring that the Internet fulfils its social and economic potential.”

Another cross-cutting facet of openness is multilingualism. If the Internet cannot accommodate a language, people who can communicate only in that language will not be able to enjoy the social and economic benefits that people who speak other languages have. Furthermore, the online contributions that could have been made by people who are linguistically blocked will be unavailable to everyone.

One of the most important characteristics of openness is inclusive governance. This means that decisions about shared principles, norms, rules, procedures and programs that shape the ways in which the Internet is used and evolves are made not just by one group, but by governments, the private sector, the technical community and civil society working collaboratively.

Finally, Internet openness involves distributed control. The Internet is not centrally managed. It depends on the voluntary participation and collaboration of many people and organizations to oversee its independent components and make the Internet work. While the various participants need to follow the Internet’s widely adopted technical protocols and standards, the distributed control arrangement allows them to organize and operate their particular parts of the Internet largely in the manner of their choosing.

From a practical standpoint, openness corresponds with the individual’s ability to use the Internet to do more things online, whether it is starting an e-business, expressing opinions, sharing knowledge and ideas, or using a map on a mobile device. Certain factors such as personal privacy, the security of commercial data, national security and fundamental values must be given due regard in determining the degree of openness that a society wishes to have. It is not the purpose of this chapter, however, to reach conclusions about how much openness or closedness there should be.

CONCLUSION

This chapter has proposed a broad definition of Internet openness. It is well known that certain technical elements of the Internet’s architecture, such as publicly available and commonly adopted data transport protocols, have had profound effects on economies and societies by virtue of their contribution to openness. By including economic, social and other elements in the definition, this chapter recognizes that Internet openness also depends on an array of non-technical factors such as affordable access, privacy rights and transparent regulations. If the implications of this definition of Internet openness can be distilled into one phrase, it is that Internet openness leads to the global free flow of data across the network.

With a working definition of Internet openness in hand, it is possible to take additional steps toward better understanding how — and how much — changes in

openness are affecting economic and social outcomes. The OECD is now taking those steps with the aim of helping policy makers to take evidence-based approaches to decisions about Internet openness.

ACKNOWLEDGEMENT

The author wishes to thank Geoff Huston, a consultant to the OECD, whose work (which can be found at www.potaroo.net/ispcol/2015-10/open.pdf) provided a basis for the more technical aspects of this chapter.

WORKS CITED

- Blumenthal, M. and D. Clark. 2001. "Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World." *ACM Transactions on Internet Technology* 1 (1): 70–109.
- Box, Sarah. 2016. *Internet Openness and Fragmentation: Toward Measuring the Economic Effects*. Global Commission on Internet Governance Paper Series No. 36. Waterloo, ON: CIGI.
- Dalberg. 2014. *Open for Business? The Economic Impact of Internet Openness*. Dalberg Global Development Advisors, March. www.dalberg.com/documents/Open_for_Business_Dalberg.pdf.
- OECD. 2008. *The Seoul Declaration for the Future of the Internet Economy*. Paris, France: OECD Publishing. www.oecd.org/sti/40839436.pdf.
- . 2013. *The OECD Privacy Framework*. Paris, France: OECD Publishing. www.oecd.org/sti/ieconomy/oecd_privacy_framework.pdf.
- . 2014. *Principles for Internet Policy Making*. Paris, France: OECD Publishing. www.oecd.org/sti/ieconomy/oecd-principles-for-internet-policy-making.pdf.
- . 2015. *Digital Security Risk Management for Economic and Social Prosperity: OECD Recommendation and Companion Document*. Paris, France: OECD Publishing. <http://dx.doi.org/10.1787/9789264245471-en>.
- Saltzer, J., D. Reed and D. Clark. 1981. "End-to-End Arguments in System Design." In *Proceedings of the Second International Conference on Distributed Computing Systems* [Paris, April 8–10]. Washington, DC: IEEE Computer Society.
- The Economist*. 2010. "The Web's New Walls: How the threats to the internet's openness can be averted." *The Economist*, September 2. www.economist.com/node/16943579.
- UN. 1948. General Assembly resolution 217 A, *Universal Declaration of Human Rights*, A/RES/217 (11) (December 10). www.un.org/en/universal-declaration-human-rights/.
- . 2012. General Assembly resolution 20.8, *The promotion, protection and enjoyment of human rights on the Internet*, A/HRC/RES/20.8 (July 16). <http://daccess-dds-ny.un.org/doc/RESOLUTION/GEN/G12/153/25/PDF/G1215325.pdf?OpenElement>.

ABOUT THE AUTHOR

Jeremy West is a senior policy analyst in the Digital Economy Policy Division of the Directorate for Science, Technology and Innovation at the Organisation for Economic Co-operation and Development (OECD). He recently led a multidisciplinary project on intellectual property in the digital age and is currently researching the economic and social effects of Internet openness. Jeremy's background is in competition law and includes experience at a law firm in Washington, DC, at the United States Department of Justice and with the New Zealand Commerce Commission. Jeremy serves on the editorial boards of the *Antitrust Law Journal* and *Oxford Competition Law*, and is a member of the State Bar of California and the District of Columbia Bar. His OECD papers have been cited by the United States Antitrust Modernization Commission, the *Financial Times* and in scholarly journals.

ABOUT CIGI

We are the Centre for International Governance Innovation: an independent, non-partisan think tank with an objective and uniquely global perspective. Our research, opinions and public voice make a difference in today's world by bringing clarity and innovative thinking to global policy making. By working across disciplines and in partnership with the best peers and experts, we are the benchmark for influential research and trusted analysis.

Our research programs focus on governance of the global economy, global security and politics, and international law in collaboration with a range of strategic partners and support from the Government of Canada, the Government of Ontario, as well as founder Jim Balsillie.

Au Centre pour l'innovation dans la gouvernance internationale (CIGI), nous formons un groupe de réflexion indépendant et non partisan qui formule des points de vue objectifs dont la portée est notamment mondiale. Nos recherches, nos avis et l'opinion publique ont des effets réels sur le monde d'aujourd'hui en apportant autant de la clarté qu'une réflexion novatrice dans l'élaboration des politiques à l'échelle internationale. En raison des travaux accomplis en collaboration et en partenariat avec des pairs et des spécialistes interdisciplinaires des plus compétents, nous sommes devenus une référence grâce à l'influence de nos recherches et à la fiabilité de nos analyses.

Nos programmes de recherche ont trait à la gouvernance dans les domaines suivants : l'économie mondiale, la sécurité et les politiques mondiales, et le droit international, et nous les exécutons avec la collaboration de nombreux partenaires stratégiques et le soutien des gouvernements du Canada et de l'Ontario ainsi que du fondateur du CIGI, Jim Balsillie.

For more information, please visit www.cigionline.org.

ABOUT CHATHAM HOUSE

Chatham House, the Royal Institute of International Affairs, is based in London. Chatham House's mission is to be a world-leading source of independent analysis, informed debate and influential ideas on how to build a prosperous and secure world for all. The institute: engages governments, the private sector, civil society and its members in open debates and confidential discussions about significant developments in international affairs; produces independent and rigorous analysis of critical global, regional and country-specific challenges and opportunities; and offers new ideas to decision-makers and -shapers on how these could best be tackled from the near- to the long-term. For more information, please visit: www.chathamhouse.org.

CIGI MASTHEAD

Executive

President	Rohinton P. Medhora
Director of Finance	Shelley Boettger
Director of the International Law Research Program	Oonagh Fitzgerald
Director of the Global Security & Politics Program	Fen Osler Hampson
Director of Human Resources	Susan Hirst
Director of the Global Economy Program	Domenico Lombardi
Chief Operating Officer and General Counsel	Aaron Shull
Director of Communications and Digital Media	Spencer Tripp

Publications

Publisher	Carol Bonnett
Senior Publications Editor	Jennifer Goyder
Publications Editor	Patricia Holmes
Publications Editor	Nicole Langlois
Publications Editor	Sharon McCartney
Publications Editor	Lynn Schellenberg
Graphic Designer	Sara Moore
Graphic Designer	Melodie Wakefield

For publications enquiries, please contact publications@cigionline.org.

Communications

For media enquiries, please contact communications@cigionline.org.



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

CHATHAM HOUSE

The Royal Institute of
International Affairs

10 St James's Square
London, England SW1Y 4LE
United Kingdom
tel +44 (0)20 7957 5700 fax +44 (0)20 7957 5710
www.chathamhouse.org