



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

Global Commission on Internet Governance

ourinternet.org

PAPER SERIES: NO. 7 — MARCH 2015

On the Nature of the Internet

Leslie Daigle



ON THE NATURE OF THE INTERNET

Leslie Daigle



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

Copyright © Leslie Daigle, 2015. This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Published by the Centre for International Governance Innovation and Chatham House.

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



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



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



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

vi	About the Global Commission on Internet Governance
vi	About the Author
1	Acronyms
1	Executive Summary
1	Introduction
2	The Technical Nature of the Internet
5	The Unvarying Characteristics That Define the Internet: The Invariants
9	Situational Challenges and Threats of Fragmentation of the Internet
15	Considering the Nature of the Internet in Policy Discussions
16	Conclusion
16	Works Cited
20	About CIGI
20	About Chatham House
20	CIGI Masthead

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 conducts and supports independent research on Internet-related dimensions of global public policy, culminating in an official commission report that will articulate concrete policy recommendations for the future of Internet governance. These recommendations will 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 will focus 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

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.

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

EXECUTIVE SUMMARY

This paper examines three aspects of the nature of the Internet: the Internet’s technology, general properties that make the Internet successful and current pressures for change. Current policy choices can, literally, make or break the Internet’s future. By understanding the Internet — primarily in terms of its key properties for success, which have been unchanged since its inception — policy makers will be empowered to make thoughtful choices in response to the pressures outlined here, as well as new matters arising.

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 paper 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.

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.

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?

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

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.

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.

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 <http://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 <http://www.routeviews.org>.

⁶ See <http://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 <http://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 <http://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 <http://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. <http://www.rfc-editor.org/rfc/rfc4033.txt>.
- Al Jazeera. 2011. “When Egypt Turned Off the Internet.” January 28. <http://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. <http://www.rfc-editor.org/rfc/rfc1436.txt>.
- Block, Alex Ben. 2014. “India Joins China, Russia, Switzerland on Privacy Watch List.” *Hollywood Reporter*, June 24. <http://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. <http://www.reuters.com/article/2014/03/19/us-brazil-internet-idUSBREA2I03O20140319>.
- Colitti, Lorenzo. 2009. “IPv6: The Acceptance Phase.” Presentation given at IETF 74, March 22–27. <http://www.isoc.org/isoc/conferences/ipv6panel/docs/colitti.pdf>.
- Crispin, M. 2003. “Internet Message Access Protocol — Version 4rev1.” RFC3501, March. <http://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. <http://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. <http://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. <http://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. <http://www.ietf.org/rfc/rfc1519.txt>.
- Goble, Gordon. 2012. "Top Ten Bad Tech Predictions," Digital Trends, November 4. <http://www.digitaltrends.com/features/top-10-bad-tech-predictions/8/>.
- Google. 2014. "Google IPv6 Traffic." <http://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. <http://www.forbes.com/sites/adamhartung/2011/01/14/why-facebook-beat-myspace/>.
- IANA. 2014. "IANA IPv4 Address Space Registry." <http://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. <http://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. <http://www.ietf.org/rfc/rfc2826.txt>.
- Internet Society. 2012. "Internet Invariants: What Really Matters." <http://www.internetsociety.org/internet-invariants-what-really-matters>.
- . 2013. "A Fine Balance: Internet Number Resource Distribution and De-centralisation." <http://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. http://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. <http://www.ietf.org/rfc/rfc5321.txt>.
- Kozlov, Vladimir. 2014. "Russia's Anti-Piracy Law to be Toughened, Producing Exec Says" *Hollywood Reporter*, October 7. <http://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. <http://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. <http://www.ietf.org/rfc/rfc3403.txt>.
- Mockapetris, P. V. 1987. "Domain Names — Concepts and Facilities." RFC1034, November. <http://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. <http://www.ietf.org/rfc/rfc4271.txt>.
- Ricknäs, M. 2008. "Sprint-Cogent Dispute Puts Small Rip in Fabric of Internet." *PC World*, October 31.

CIGI PUBLICATIONS

ADVANCING POLICY IDEAS AND DEBATE



The Regime Complex for Managing Global Cyber Activities

GCIG Paper No. 1
Joseph S. Nye, Jr.

When trying to understand cyber governance, it is important to remember how new cyberspace is. Advances in technology have, so far, outstripped the ability of institutions of governance to respond. This paper concludes that predicting the future of the normative structures that will govern cyberspace is difficult.



Legal Interoperability as a Tool for Combatting Fragmentation

GCIG Paper No. 4
Rolf H. Weber

The recently developed term “legal interoperability” addresses the process of making legal rules cooperate across jurisdictions. It can facilitate global communication, reduce costs in cross-border business and drive innovation, thereby creating a level playing field for the next generation of technologies and cultural exchange.



Tipping the Scale: An Analysis of Global Swing States in the Internet Governance Debate

GCIG Paper No. 2
Tim Maurer and Robert Morgus

This paper offers an analysis of the global swing states in the Internet governance debate and provides a road map for future in-depth studies.



Innovations in Global Governance: Toward a Distributed Internet Governance Ecosystem

GCIG Paper No. 5
Stefaan G. Verhulst, Beth S. Noveck, Jillian Raines and Antony Declercq

The growth and globalization of the Internet over the past 40 years has been nothing short of remarkable. Figuring out how to evolve the Internet’s governance in ways that are effective and legitimate is essential to ensure its continued potential.



Legal Mechanisms for Governing the Transition of Key Domain Name Functions to the Global Multi-stakeholder Community

GCIG Paper No. 3
Aaron Shull, Paul Twomey and Christopher S. Yoo

Under the existing contractual arrangement, the Internet Corporation for Assigned Names and Numbers (ICANN) has been accountable to the US government for the performance of these functions. However, if the US government is no longer party to this agreement, then to whom should ICANN be accountable?



The Impact of the Dark Web on Internet Governance and Cyber Security

GCIG Paper No. 6
Michael Chertoff and Toby Simon

The deep Web has the potential to host an increasingly high number of malicious services and activities. The global multi-stakeholder community needs to consider its impact while discussing the future of Internet governance.

ORGANIZED CHAOS

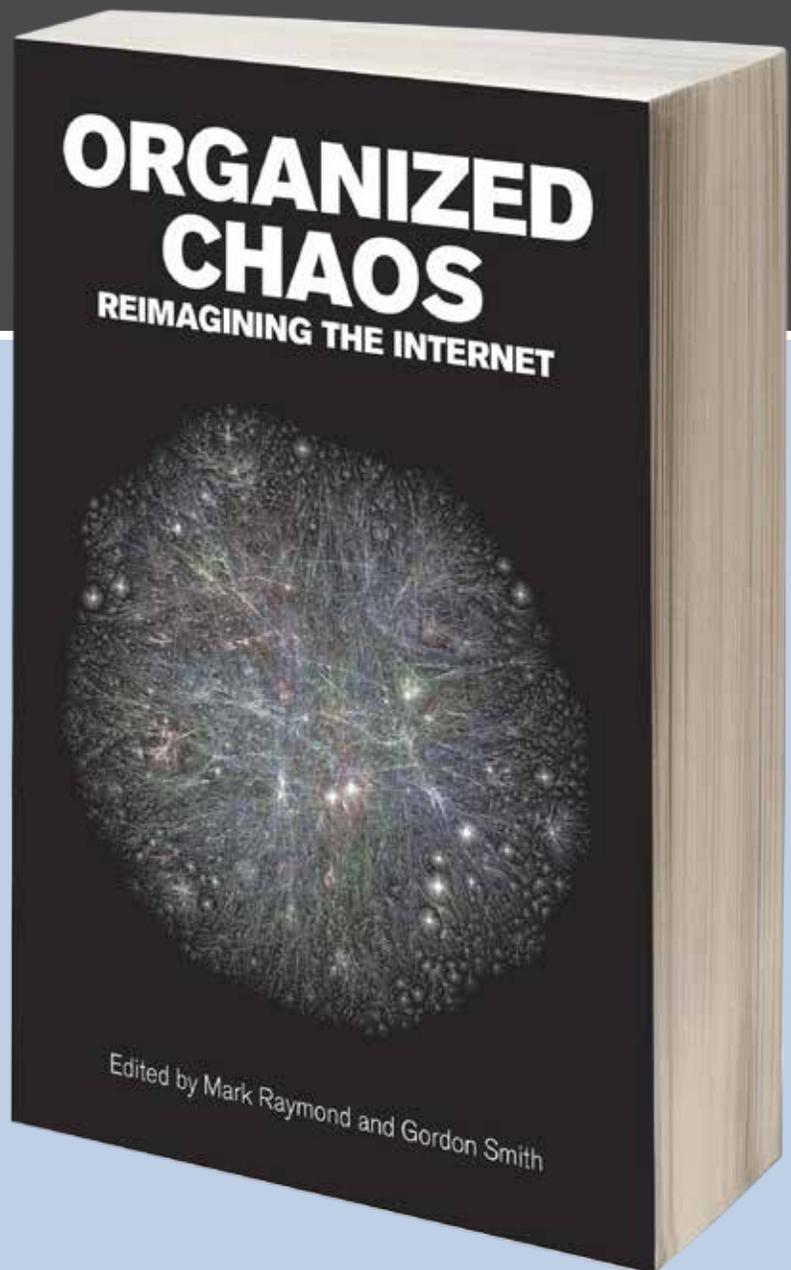
REIMAGINING THE INTERNET

EDITED BY

MARK RAYMOND AND
GORDON SMITH

Leading experts address a range of pressing challenges, including cyber security issues and civil society hacktivism by groups such as Anonymous, and consider the international political implications of some of the most likely Internet governance scenarios in the 2015-2020 time frame. Together, the chapters in this volume provide a clear sense of the critical problems facing efforts to update and redefine Internet governance, the appropriate modalities for doing so, and the costs and benefits associated with the most plausible outcomes. This foundation provides the basis for the development of the research-based, high-level strategic vision required to successfully navigate a complex, shifting and uncertain governance environment.

Price: CDN\$25.00
200 Pages, Trade Paperback
ISBN 978-1-928096-04-7



Centre for International Governance Innovation

Single copy orders: cigionline.org/bookstore
Available in paperback and ebook form.

ABOUT CIGI

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

CIGI's current research programs focus on three themes: the global economy; global security & politics; and international law.

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

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

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

Managing Editor, Publications	Carol Bonnett
Publications Editor	Jennifer Goyder
Publications Editor	Vivian Moser
Publications Editor	Patricia Holmes
Publications Editor	Nicole Langlois
Graphic Designer	Melodie Wakefield
Graphic Designer	Sara Moore

EXECUTIVE

President	Rohinton Medhora
Vice President of Programs	David Dewitt
Vice President of Public Affairs	Fred Kuntz
Vice President of Finance	Mark Menard

COMMUNICATIONS

Communications Manager	Tammy Bender	tbender@cigionline.org (1 519 885 2444 x 7356)
-------------------------------	--------------	--



67 Erb Street West
Waterloo, Ontario N2L 6C2
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