

Research Volume Five

Global Commission on Internet Governance

Cyber Security in a Volatile World



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

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

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

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

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

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



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

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



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

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

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

TABLE OF CONTENTS

About the Global Commission on Internet Governance iv

Preface v
Carl Bildt

Introduction: Security as a Precursor to Internet Freedom and Commerce 1
Laura DeNardis

Chapter One: Global Cyberspace Is Safer than You Think: Real Trends in Cybercrime 5
Eric Jardine

Chapter Two: The Impact of the Dark Web on Internet Governance and Cyber Security 29
Michael Chertoff and Toby Simon

Chapter Three: The Dark Web Dilemma: Tor, Anonymity and Online Policing 37
Eric Jardine

Chapter Four: The Tor Dark Net 51
Gareth Owen and Nick Savage

Chapter Five: Connected Choices: How the Internet Is Challenging Sovereign Decisions 63
Melissa E. Hathaway

Chapter Six: Increasing Internet Connectivity While Combatting Cybercrime: Ghana as a Case Study 77
Caroline Baylon and Albert Antwi-Boasiako

Chapter Seven: Critical Infrastructure and the Internet of Things 93
Toby Simon

Chapter Eight: Combatting Cyber Threats: CSIRTs and Fostering International Cooperation on Cyber Security 105
Samantha Bradshaw

Chapter Nine: Toward a Social Compact for Digital Privacy and Security 121
Statement by the Global Commission on Internet Governance

About CIGI. 133

About Chatham House 133

CIGI Masthead. 133

ABOUT THE GLOBAL COMMISSION ON INTERNET GOVERNANCE

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

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

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

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

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

www.ourinternet.org

PREFACE

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

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

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

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

Carl Bildt

Chair, GCIG

November 2016

**INTRODUCTION:
SECURITY AS A PRECURSOR TO INTERNET FREEDOM AND COMMERCE**
Laura DeNardis

Copyright © 2017 by Laura DeNardis

INTRODUCTION

Security as a Precursor to Internet Freedom and Commerce

The global digital economy, democracy and the public sphere now completely depend upon the stability and security of cyberspace. Encryption technologies are necessary to protect data privacy, authenticate websites and secure online transactions. Security problems such as consumer data breaches and denial-of-service attacks disrupt the digital economy and the public sphere. They also can have chilling effects on speech and online behaviour. As everyday physical objects from cars to home appliances increasingly become Internet-connected, human safety in the real world also depends upon cyber security. Trust in digital infrastructure is now necessary for the capacity to communicate, access knowledge, use one's banking system, drive a car and buy products through an online commerce site such as Amazon. Democracy also depends upon cyber security, considering the stunning admission by United States intelligence agencies about Russia's influence campaign, probing of voter rolls and hacking of Democratic National Committee emails during the 2016 presidential campaign. Cyber security is one of the great human rights issues of our time.

Cyber security is not only an issue for "Internet users" but for all citizens. Even someone who has never been online is directly affected when a retail company they frequent (for example, Target or Home Depot) experiences a massive consumer data breach, when their television potentially becomes a surveillance tool or when they are denied medical care because of a ransomware attack that cryptographically locks medical records and otherwise disables health care provider systems. All people and all societies are now directly affected by the security of digital systems. Trust in these systems is a precursor to basic functioning in the modern world.

Quantitative studies suggest that trust in cyberspace has direct implications for user behaviour, commerce and freedom. To help support the work of the Global Commission on Internet Governance (GCIG), the Centre for International Governance Innovation (CIGI) undertook, over the course of multiple years (2014, 2016, 2017), massive polls of Internet users around the globe to gauge public trust and understanding in a variety of cyber security and Internet governance areas. For example, the 2017 CIGI-Ipsos Global Survey on Internet Security and Trust polled 24,225 Internet users across 24 countries: Australia, Brazil, Canada, China, Egypt, France, Germany, Hong Kong (China), India, Indonesia, Italy, Japan, Kenya, Mexico, Nigeria, Pakistan, Poland, Republic of Korea, South Africa, Sweden, Tunisia,

Turkey, the United Kingdom and the United States.¹ The 2017 poll indicated that a majority of global respondents were more concerned about their online privacy relative to the previous year, and that the top concern among those worried about their privacy was cybercrime (82 percent), followed by Internet privacy companies (74 percent) and governments (65 percent) (see Table 1).

This concern is perhaps explainable by respondents' awareness of high-profile consumer-data breaches of major retail, financial and insurance companies, but also their recognition of government surveillance of citizens and corporate data-gathering practices that collect and share personal information to support business models based on targeted online advertising.

This research volume sets out to quantitatively and qualitatively examine the state of global cyber security and address what can be done differently to improve security, stability and trust online. Even in the midst of rising public awareness and concern about cyber security breaches and digital crime, research scholar Eric Jardine quantitatively examines, in *Global Cyberspace Is Safer than You Think: Real Trends in Cybercrime* (2015a), the question of the actual trends in cybercrime. He argues that the level of security in cyberspace is, in reality, far better than depictions of the level of security. Cybercrime statistics are usually depicted in absolute numbers and present year-over-year statistics that fail to normalize numbers to account for the growth in cyberspace. These absolute numbers, Jardine demonstrates, are far worse than the normalized numbers.

One challenging area of cyber security policy intervention is the so-called dark Web. The term "dark Web" generally refers to sites or collections of sites on the Internet that are intentionally hidden, not accessible via standard browsers, and that rely on strong anonymization and encryption technologies to enable secrecy. It generally connotes an area in which a great deal of malicious and illegal transactions occur, such as human trafficking and illegal drugs and arms sales. In other cases, the same technologies are used to enable journalist activities or dissident communication in authoritarian political environments. Three papers address various tensions arising in the dark Web. *The Impact of the Dark Web on Internet Governance Cyber Security* (2015), by Michael Chertoff and Toby Simon, provides an overview of the implications of the dark Web and how it should fit into the broader Internet governance milieu. Eric Jardine addresses the policy dilemmas around anonymity-granting systems such as The Onion Router (TOR) in *The Dark Web Dilemma: Tor, Anonymity and Online Policing* (2015b). Researchers Gareth Owen and Nick Savage, in *The Tor Dark Net* (2015), present the results of their collection and analysis of six months of data about Tor sites, finding that

¹ Additional background materials and results of the 2017 CIGI-Ipsos Global Survey on Internet Security Trust are available at <https://www.cigionline.org/internet-survey>.

Table 1: Top Privacy Concerns

	How Concerned?		
	A Great Deal	Somewhat	A Great Deal or Somewhat (Total)
Cybercriminals	56%	26%	82%
Internet companies	35%	39%	74%
Other Internet users	29%	38%	67%
Your government	32%	33%	65%
Companies in general	22%	40%	62%
Foreign governments	26%	35%	61%
Employers	17%	32%	49%

Data source: 2017 CIGI-Ipsos Global Survey on Internet Security and Trust.

the majority were used to carry out criminal activities, in particular involving drug marketplaces and child abuse images.

The role of sovereign nation-states in addressing cyber security is in a state of flux. On the one hand, these states have an interest in preserving the security of critical infrastructure. On the other hand, governments have an interest in weak security that allows them to more easily carry out law enforcement and intelligence activities and engage in geopolitically motivated cyber conflict. Commissioner and cyber security expert Melissa Hathaway tackles the competing interests and stakes over control of cyberspace in *Connected Choices: How the Internet Is Challenging Sovereign Decisions* (2015). As a more specific case study, Caroline Baylon and Albert Antwi-Boasiako examine interlinkages between Internet development and cyber security in *Increasing Internet Connectivity While Combatting Cybercrime: Ghana as a Case Study* (2016).

Many discussions of cyber security, whether in private industry fora, the policy-making realm or academic circles, are still geared toward Internet systems designed for the exchange of communication among people. Yet, already, there are more things than people connected to the Internet. The terms “Internet of Things” (IoT) and “cyber physical systems” address the growing realm of the Internet in which objects, industrial control systems, medical devices, wearable technologies and other material devices are digitally interconnected and tied to the public Internet. IoT cyber security will be one of the great public policy issues of the current generation. GCIG Commissioner Bobby Simon addresses this topic in his paper *Critical Infrastructure and the Internet of Things* (2017).

As cyber threats become more complicated, the role of institutions such as computer security incident response teams (CSIRTs) become more important. However, there is often not trust and information sharing among various CSIRTs. Security researcher Samantha Bradshaw, in her paper *Combatting Cyber Threats: CSIRTs and Fostering International Cooperation on Cyber Security* (2015), helps to

explain the role of CSIRTs in the broader “cyber regime complex” and explains some of the factors leading to information sharing and trust problems.

Not surprisingly, decisions about cyber security often exist at the nexus of competing values, such as national security versus individual privacy and digital commerce. For example, strong encryption is necessary for authenticating and securing financial and other commercial transactions online, and is also necessary for protecting individual privacy. But strong encryption comes into tension with the need for law enforcement and intelligence agencies to combat terrorism, cybercrime, identity theft and piracy online. The final chapter of this volume is a formal statement by the Global Commission on Internet Governance addressing this tension between privacy and security. *Toward a Social Compact for Digital Privacy and Security* (2015), published herein in its entirety, “calls on the global community to build a new social compact between citizens and their elected representatives, the judiciary, law enforcement and intelligence agencies, business, civil society and the Internet technical community, with the goal of restoring trust and enhancing confidence in the Internet” (GCIG 2015, 1).

The Commission’s Social Compact chapter is a strong statement both for privacy and the rule of law, suggesting ways in which these values are not mutually exclusive, but rather mutually reinforcing. The Commission recognizes privacy and personal data protection as fundamental human rights and develops reasonable parameters for government surveillance by law enforcement and intelligence agencies. Specifically, surveillance should be “openly specified in advance, authorized by law and consistent with the principles of necessity and proportionality” (GCIG 2015, 2). In the contemporary context of conflicts between private industry and governments on the strength of and limitations on surveillance, the Commission takes a firm stand that governments should not create or require back doors to access encrypted data that would, in effect, weaken Internet security.

Taken together, the Commission's call for a new social compact for digital privacy and security, along with the research chapters include in this volume, identify a wide swath of contemporary cyber security challenges and policy recommendations for providing the necessary cyber stability and security to sustain the digital economy and protect the day-to-day dependencies on cyberspace.

WORKS CITED

- Baylon, Caroline and Albert Antwi-Boasiako. 2016. *Increasing Internet Connectivity While Combatting Cybercrime: Ghana as a Case Study*. GCIG Paper Series No. 44. Waterloo, ON: CIGI. www.ourinternet.org/research/increasing-internet-connectivity-while-combatting-cybercrime-ghana-case-study.
- Bradshaw, Samantha. 2015. *Combatting Cyber Threats: CSIRTs and Fostering International Cooperation on Cyber Security*. GCIG Paper Series No. 23. Waterloo, ON: CIGI. www.ourinternet.org/research/combating-cyber-threats-csirts-and-fostering-international-cooperation-cybersecurity.
- Chertoff, Michael and Toby Simon. 2015. *The Impact of the Dark Web on Internet Governance and Cyber Security*. GCIG Paper Series No. 6. Waterloo, ON: CIGI. www.ourinternet.org/research/impact-dark-web-internet-governance-and-cyber-security.
- Global Commission on Internet Governance. 2015. *Toward a Social Compact for Digital Privacy and Security*. Waterloo, ON: CIGI. www.ourinternet.org/research/toward-social-compact-digital-privacy-and-security.
- Hathaway, Melissa E. 2015. *Connected Choices: How the Internet Is Challenging Sovereign Decisions*. GCIG Paper Series No. 11. Waterloo, ON: CIGI. www.ourinternet.org/research/connected-choices-how-internet-challenging-sovereign-decisions.
- Jardine, Eric. 2015a. *Global Cyberspace Is Safer than You Think: Real Trends in Cybercrime*. GCIG Paper Series No. 16. Waterloo, ON: CIGI. www.ourinternet.org/research/global-cyberspace-safer-you-think-real-trends-cybercrime.
- . 2015b. *The Dark Web Dilemma: Tor, Anonymity and Online Policing*. GCIG Paper Series No. 21. Waterloo, ON: CIGI. www.ourinternet.org/sites/default/files/publications/n21.pdf.
- Owen, Gareth and Nick Savage. 2015. *The Tor Dark Net*. GCIG Paper Series No. 20. Waterloo, ON: CIGI. www.ourinternet.org/research/tor-dark-net.
- Simon, Toby. 2017. *Critical Infrastructure and the Internet of Things*. GCIG Paper Series No. 46. Waterloo, ON: CIGI. www.cigionline.org/publications/critical-infrastructure-and-internet-things-0.

ABOUT THE AUTHOR

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

**CHAPTER ONE:
GLOBAL CYBERSPACE IS SAFER THAN YOU THINK:
REAL TRENDS IN CYBERCRIME**

Eric Jardine

Copyright © 2014 by the Centre for International Governance Innovation
and the Royal Institute of International Affairs

In politics, what begins in fear usually ends in folly.

– Samuel Taylor Coleridge

INTRODUCTION

Recent media coverage has been chock full of high-profile accounts of cybercrime. Hacks, data breaches, destruction of property and the theft of personal information seems to be rampant. In February 2014, eBay's online system was breached after some of its employees' credentials were stolen, leading to the compromise of some 145 million account holders (Finkle, Chatterjee and Maan 2014). In July, the American bank JPMorgan Chase was hacked, with online bandits making off with account information on approximately 76 million households and some eight million small businesses (Silver-Greenberg, Goldstein and Perloth 2014). In November, Sony Pictures was subject to a sophisticated cyber attack, causing massive physical damage to its computer systems and exposing sensitive emails regarding pay disparities and personal relationships. In December 2014, Sony estimated that the remediation and investigation costs of the hack could enter into the \$100 million¹ range (Richwine 2014). What is more, these are just a few of the publicly known breaches.

As the Internet comes to underwrite more and more of our daily life, the vectors of attack for cybercriminals, hackers and state officials multiply, the total number of cyber attacks grows year over year and the potential damage from cyber attacks increases. Governments, corporations and individuals have prudently responded to these trends by stepping up their cyber defences. Shortly after the Sony Pictures hacks, for example, the United States and the United Kingdom announced a series of "cyber war games" to prepare their government agencies for the potential of broad-based cyber attacks on critical infrastructure, including the banking and financial sector (BBC News 2015). Over 60 percent of businesses' representatives surveyed in a recent Gandalf Group C-Suite study have responded to the perception of a deteriorating cyber security environment by increasing their information technology (IT) security budgets (Gandalf Group 2014). Likewise, a recent CIGI-Ipsos poll surveying over 23,000 respondents in 24 countries found that 64 percent of respondents were more worried about their online privacy compared to one year ago and 78 percent of respondents were concerned about criminal hackers stealing their banking information. An additional 77 percent of respondents were concerned that online criminals would steal their private messages and photos. Indicating the behavioural changes that people have undertaken in response to perceptions of the poor security of cyberspace, the survey also found that compared to one year ago, some 43 percent of respondents now avoid certain Internet sites and web applications,

about 39 percent change their passwords regularly and roughly 10 percent actually use the Internet less often (CIGI-Ipsos 2014).

Clearly, the proliferation of cybercrime and the media's coverage of high-profile hacks have generated a severely negative perception of the security of cyberspace and caused governments, businesses and individual citizens to take additional steps to protect themselves online. The problem is that the existing picture of the security of cyberspace is misleading. Currently, statistics on cybercrime are, as far as I am aware, always expressed in either absolute (1,000 attacks per year) or year-over-year (2013 had 46 percent more cyber attacks than 2012) terms.² The difficulty with this expression of the numbers is that it gives an inaccurate picture of the actual trends in cybercrime over time, and thus a false impression of the actual security of cyberspace. To state the obvious (but perhaps not well understood), the occurrence of cybercrime is inevitably related to the size of the Internet. Since cyberspace is, in a number of ways, expanding at an exponential rate, it is reasonable to expect that the *absolute* number of cyber attacks will also increase simply because the Internet ecosystem is getting bigger and not necessarily because the situation is growing worse. These observations raise two questions: What is the actual trend in cyber security? And is cyberspace becoming less safe, safer or staying roughly the same over time?

In order to provide an accurate picture of the security of cyberspace, all indicators of cybercrime need to be normalized around data that captures the growing size of the Internet.³ An example to help clarify the importance of normalizing (or, essentially, expressing numbers as a proportion of a population) data on cybercrime around the size of the Internet is as follows: Imagine there is a town of 1,000 people with 100 violent crimes a year. Now imagine that there is a city with 100,000 people with 1,000 violent crimes per year. When normalizing the crime statistics for these two hypothetical population centres, it is found that the town has a violent crime rate of 0.1, while the city has a violent crime rate of 0.01. In other words, even though the city has as many violent crimes as the entire population of the town, a person's chance of being subject to a violent

2 The two exceptions involve spam and phishing emails, often expressed as a percentage of all emails sent. There is no clear rationale given for why cybercrime statistics are expressed in absolute or year-over-year terms. One potential reason is that, as shown in this chapter, the numbers tend to be more severe and point to a worse situation. Since most collectors of cybercrime data are private, for-profit companies, a cynic could conclude that the companies present data in a specific way to help them sell product. I have no proof at all of this interpretation. It is merely one potential explanation.

3 In this chapter, the terms Internet and cyberspace are used synonymously. The Internet usually refers to the physical structure of the network, while cyberspace is the larger, over-the-top portion of the Web involving things such as apps. Both terms herein mean cyberspace and both are used in the chapter to mean the same thing in the interest of readability.

1 All currency is in US dollars.

crime in the city is only 1 in 100, while the chance of being the victim of a violent crime in the town is 1 in 10.

In the case of the global Internet, the occurrence of cybercrime can only be meaningfully normalized around figures that capture the full width and breadth of cyberspace. Cyber attacks in one country can originate in any other country on the planet that has an Internet connection. Normalizing crime statistics around national-level data, therefore, gives a partial and highly skewed glimpse at real trends in the occurrence and cost of cybercrime.

Taking data on the size of the Internet and normalizing various cybercrime indicators around these figures from 2008 to the end of 2014, the security of cyberspace is better than one would think from looking at just the absolute numbers often presented in the media and in IT security reports. Over 30 comparisons of the absolute (1,000 attacks) and normalized (0.15 attacks per 1,000 Internet users) numbers bear out this claim.

When the normalized indicators of cybercrime are compared to the absolute numbers that are usually used to discuss the level of security in cyberspace, one of three misrepresentations occurs:

- the absolute numbers indicate the situation is getting worse when the normalized numbers say it is getting better (as in the case of new vulnerabilities, zero-day vulnerabilities, browser vulnerabilities, mobile vulnerabilities, post-breach response costs and notification costs);
- both the absolute and the normalized numbers say the situation is worsening, but the absolute numbers say it is growing worse at a faster rate than the normalized numbers (as in the case of detection and escalation costs, when the full sample is considered); or
- both the absolute and the normalized numbers say the situation is improving, but the absolute numbers indicate a slower rate of improvement than the normalized numbers (as in the case of malicious web domains, botnets, web-based attacks since 2012, average per capita data breach costs, organizational costs due to data breaches, detection and escalation costs from 2010 to 2013 or lost business costs).

In short, when the number of cyber attack vectors, the number of cyber attacks and the amount of damage caused by cybercrime are expressed as a proportion of the size of the Internet, each of the normalized numbers point to the idea that the security of cyberspace is better than is suggested by the un-normalized or absolute numbers. As a result, the security of cyberspace is likely better than is commonly perceived by the general public, private companies and state officials.

A realistic understanding of the level of security in cyberspace is important because an unnecessarily negative image of the situation can lead to radical policy responses that could easily produce more harm than good. If online crime is rampant, then restricting online activity might be warranted, likely to the ultimate detriment of cultural expression, commerce and innovation. If, on the other hand, cyberspace security is relatively good, then current policies could be sufficient and things can go on more or less as they do now. In any case, a more realistic impression of the security of cyberspace provides a better foundation for cyber security policy.

The chapter first discusses how to conceptualize the size of cyberspace and details the data that is used herein to measure this concept. It then provides a three-part framework for thinking about the security of cyberspace and details the measures used to operationalize each part of the framework. The next three sections examine the normalized trends in each of these areas and compares them to the trends in the absolute numbers. The chapter concludes with policy recommendations based on the finding that cyberspace security is better than what is indicated when looking at only the absolute numbers and is actually, in many cases, getting better rather than worse.⁴

THE SIZE OF CYBERSPACE

The cyberspace ecosystem is built upon the physical infrastructure of the Internet and is basically composed of users, points of online interaction (websites, for instance) and the volume of activity that occurs online. The online ecosystem gets larger as the number of users, points of interaction and volume of activity increases. This section lays out a three-part framework for understanding the scope, size, width and breadth of cyberspace. Cyberspace is essentially an amalgamation of the number of users (people and devices, etc.), the number of points of interaction (websites and domains, etc.) and the activity linking these broad categories (data flows and commerce, etc.).⁵

⁴ Readers interested solely in the difference between absolute and normalized numbers, rather than the method of measuring these numbers, can skip ahead to the section “Trends in the Vectors of Attack: Vulnerabilities and Malicious Sites.”

⁵ Studying cyberspace from an empirical perspective involves a bit of irony. While we live in the age of big data, where nearly everything a person does online is tracked and recorded, most of this information is proprietary and fragmented among numerous private actors. The result is that it is not easy to get a clear picture of either the size of the Internet or the occurrence of cybercrime. Data, therefore, have to be drawn from multiple sources and often estimates have to be used in place of actual figures. As a disclaimer: all the data used in this chapter presents at best a partial view of the actual ins and outs of cyberspace. Despite the fact that many of the sources consulted lay out their data collection procedures, it is not clear how random of a sample of Internet activity the data actually depicts, and so extrapolating from these findings to the entirety of cyberspace can only be done with great care.

Table 1: The Size of Cyberspace

	Minimum	Maximum	Mean	Standard Deviation
Internet Users	1,562,067,594	2,925,249,355	2,252,889,661	500,996,210
Email Users	1,300,000,000	2,504,000,000	1,951,333,333	514,583,586
Active Mobile Broadband Accounts	422,000,000	2,693,000,000	1,318,000,000	808,928,097
Number of Smartphones	139,290,000	1,244,890,000	567,862,857	419,380,858
Number of Domains	177,000,000	288,000,000	230,042,857	41,667,488
Number of Websites	172,338,726	968,882,453	471,754,976	307,845,943
Volume of Data Flows (Gigabytes)	1.2209x10 ¹¹	7.6685x10 ¹¹	4.10154x10 ¹¹	2.46421x10 ¹¹
Volume of Mobile Data (Gigabytes)	396,816,000	42,336,000,000	13,020,825,714	15,811,807,798
Number of Google Searches	637,200,000,000	2,161,530,000,000	1,538,311,571,429	5.83699x10 ¹¹
Internet's contribution to GDP (Boston Consulting Group)	1.92x10 ¹²	2.45x10 ¹²	2.19207x10 ¹²	2.18547x10 ¹¹
Internet's contribution to GDP (McKinsey & Company)	1.42x10 ¹²	1.72x10 ¹²	1.57879x10 ¹²	1.25132x10 ¹¹

The basic point is that the ecosystem of cyberspace is big and getting a lot bigger at a fairly rapid pace. This growth is akin to the growth of a population in a city or country, in the sense that a fixed amount of crime and a growing population will result in a lower crime rate or a better chance that one will not be subject to a crime.

As detailed below, data was collected from a variety of sources on the following variables for the concept of Internet users:

- the number of Internet users;
- the number of email users;
- the number of active mobile broadband subscriptions; and
- the number of smartphones sold to end-users.

The following data was collected on the concept of points of online interaction:

- the number of domains; and
- the number of websites.

And on the volume of online activity:

- the volume of total data flows;
- the volume of mobile data flows;
- the annual number of Google searches; and
- the Internet's contribution to GDP.

Table 1 provides some basic summary statistics for the data capturing the size of cyberspace.

Internet Users

The number of Internet users is a good measure of the size of cyberspace because it shows the actual number of people that are a part of the “network of networks.” In this sense, it is akin to the number of people in a city or country. It is also a good proxy for the number of devices online, although this number surpassed that of humans on the network around 2008 (Evans 2011). Data on the number of Internet users from 2008 to the end of 2014 was taken from the website Internet Live Stats, which provides real-time statistics on various indicators of the size of the Internet (Internet Live Stats 2015a).

Email is one of the most basic uses of the Internet. The number of email users online is a good measure of the size of the active population base of the online ecosystem because it captures not just the number of people who have Web access (as done via Internet users statistics), but also the number of people who actually use the Internet as a part of their daily lives. Email users, therefore, are an active subset of all Internet users. In 2014, for example, there were 421,249,355 more Internet users than email users for that year. Data on email users from 2008 to 2012 was taken from a data aggregation blog called Royal Pingdom, which is operated by the website monitoring company Pingdom (Royal Pingdom 2009; 2010; 2011; 2012; 2013). Data for email users for 2013 and 2014 were taken from a Radicati Group (2013) study of the email market.

Increasingly, people access the Internet via a mobile platform rather than a traditional desktop computer. In January 2014, mobile usage surpassed desktop usage in the United States for the first time (O'Toole 2014). The trend is even more pronounced in the developing world, where Internet access has expanded primarily by skipping the fixed access/desktop stage and moving directly into

the mobile/wireless broadband stage. Active mobile broadband subscriptions are a measure of individuals who access the Internet via a mobile device, such as a smartphone or tablet. They are a smaller, yet rapidly growing, subset of all Internet users. Data on active mobile broadband subscriptions is taken from the International Telecommunication Union's statistics (International Telecommunication Union 2015).

One user can operate multiple devices online (Evans 2011). Each device can potentially be subject to a cybercrime, meaning one person can be targeted multiple times even if one device is only targeted once. Data on the number of smartphones sold to end-users per year is used as a rough proxy for the number of devices online. The number is far, far smaller than the actual number of devices connected to the Web at any one time, but it is likely indicative of the growing trend in connected devices. Data on the number of smartphones sold to end-users is taken from Statista (2015).

Points of Online Interaction

Domains give a good sense of the size of the online ecosystem, as they are a key point of interaction with users. Internet domains include generic top-level domains (such as .com or .net) and country top-level domains (such as .ca and .uk). All domains are registered with the Domain Name System (DNS), which ensures that each domain is globally unique and that when you type in a web address you are taken to the correct website. Data on the number of domains from 2008 to 2014 is taken from Verisign's *Domain Name Industry Briefs* (2008; 2009; 2010; 2011; 2012; 2013; 2014).

The number of websites online is again a good measure of the number of points of interaction online and so a good measure of the size of the Internet ecosystem. There is significant overlap between websites and domains, although the number of websites is larger because one website can have multiple subsidiary pages and because not all websites are actually a part of the DNS. In 2014, the number of websites was 680,882,453 higher than the number of domains. Data on websites is taken from Internet Live Stats (2015b) for the period 2008 to 2014.

Volume of Online Activity

The Internet is essentially a hyperefficient way to send and receive data. Statistics on the volume of data that traverses the Internet, therefore, is a useful measure of how busy the Internet ecosystem is year over year. The Internet is composed of a number of privately run networks that interconnect to provide the system with a global reach (Woodcock and Adhikari 2011). Each network maintains its own records, and piecing together exactly how much data flows globally is extremely difficult. As such, any figure for the size of global data flows is only an estimate. For this chapter, data on the volume of Internet traffic from

2008 to 2013 was gathered from the "2009 Cisco Visual Networking Index: Forecast and Methodology, 2008–2013" and data on 2014 was taken from the 2010 iteration of this white paper (Cisco Systems 2009; 2010). The data taken from these reports are Cisco Systems' estimates of global Internet traffic flows. Despite the best efforts of Cisco Systems engineers, the data probably under-represent the true size of data flows across the Internet. They also fail to distinguish between the types of data flows (that is, streaming video versus emails and website visits), which could affect the appropriateness of normalizing cybercrime numbers around this metric.

Mobile traffic is a smaller, but rapidly growing, subset of all Internet traffic. Mobile traffic gives a rather obvious impression of how much people are using cyberspace via a mobile device. Mobile operating systems and security systems are distinct from traditional desktop-style systems, with their own weaknesses and vulnerabilities. The volume of mobile traffic shows how much mobile devices are used to access the Internet and, correspondingly, how likely they are to be the subject of a cybercrime. Data of mobile traffic is also taken from Cisco's two forecasting reports.

The Internet is also, as it is colloquially known, an "information superhighway." Another measure of the activity that occurs on the Internet, therefore, is the number of search engine queries per year. Data on the annual number of Google searches was used as a measure for Internet search queries (Statistics Brain 2015). Globally, Google Chrome is also the largest web browser in every region of the world (StatsCounter 2015). These trends suggest that Google searches are a good proxy for the occurrence of Internet-based searches more generally.

The Internet is becoming increasingly integrated into every aspect of society. One of the most meaningful (or at least most measureable) effects of this growing integration and importance is the Internet's share of global GDP. Currently, no comprehensive time series data exists for this measure. To operationalize the Internet's contribution to global GDP, two separate estimates on the Internet's contribution to various nations' GDP are used here. First is a McKinsey & Company estimate on the contribution of the Internet to the economy of 13 large nations in 2009.⁶ Together, these 13 nations make up some 70 percent of the world's GDP. Although the Internet's contribution to global GDP is likely larger than outlined in the McKinsey & Company study, the findings are fairly indicative of the Internet's general effect on global GDP. The second measure for the size of the global Internet economy is from a Boston Consulting Group study that looks at the Internet's contribution to GDP in Group of Twenty (G20) nations in the year 2010

6 The countries included in the McKinsey study are Brazil, Canada, China, France, Germany, India, Italy, Japan, the Russian Federation, South Korea, Sweden, the United Kingdom and the United States (Pélessié du Rausas et al. 2011).

(Dean et al. 2012). Together, the G20 makes up around 70 percent of the world's population and close to 90 percent of Global GDP (Griffith-Jones, Helleiner and Woods 2010, 25). Again, the Boston Consulting Group's study provides a partial, but still strongly indicative, picture of the Internet's contribution to global GDP. On average, and this is important to note for the later analysis, the Boston Consulting Group's 2010 estimates of the Internet's contribution to the global economy are, as one would expect, larger than the McKinsey & Company's estimates for the size of the Internet's contribution in 2009. This is in line with the rather intuitive idea that the Internet's contribution to the global economy is becoming proportionately more important over time. The Boston Consulting Group's figures are also more representative of the global contribution of the Internet because they include more countries. As such, even though the McKinsey & Company and the Boston Consulting Group estimates point to similar patterns vis-à-vis the absolute numbers, this chapter relies on the more inclusive estimates of the latter in the analysis below.

One additional assumption involving the GDP numbers needs to be laid bare. Both studies provide only a static snapshot of the Internet's contribution to global GDP, one in 2009 and one in 2010. In using these data in the comparisons below, it is assumed that the Internet's proportional contribution to each country's GDP remains constant, so if, as in the case of Sweden in the McKinsey & Company study, the Internet contributed 6.3 percent to the country's GDP in 2009, it is assumed that it also contributed 6.3 percent in 2008 and will only contribute that amount moving forward from 2009 into 2013. Since the Internet and Internet-enabled platforms are becoming increasingly common in business, industry and commerce, this assumption likely works against the real world trend of the Internet expanding in its importance to the economy year over year. The assumption is necessary, however, to get enough data in normalized cybercrime trends against an indicator of the economic size and importance of the Internet. This assumption will effectively under-represent the growing size of the Internet economy and thus shrink the denominator in the normalization of cybercrime statistics below. The assumption (although needed) will paint a picture of the security of cyberspace that is likely worse than what actually exists.

THE SECURITY OF CYBERSPACE: VECTORS, OCCURRENCE AND DAMAGE

The security of cyberspace can be conceptualized best from a user's perspective, broadly defined. A secure cyberspace is one in which a user can make use of the Internet without an unreasonable fear of suffering a high cost, with cost being defined in some combination of reputational,

monetary and rights violations terms. An insecure cyberspace environment is the opposite, or basically one in which using the Internet is likely to impose a large cost upon the user. This section outlines how to operationalize the level of security in cyberspace by looking at the available vectors for attack, the occurrence of online cyber attacks and the costs of successful attacks. Together, these three categories give a sense of how insecure cyberspace is for an individual user.

Many aspects of the security of cyberspace are worsening over time, but many others are actually remaining fairly static year over year. In the odd case, a given indicator is actually improving. These measures of the insecurity of cyberspace are akin to the crime rate in a city or country. If they are increasingly slower than the population, staying the same size as the population grows, or improving as the population increases, the common result is an improved crime rate.

This conceptualization of the security of cyberspace can be expressed as a function of three factors:

- the vectors available for cyber attack;
- the occurrence of cyber attacks; and
- the damage caused by successful cyber attacks.

Together, these three factors determine how secure cyberspace is for an individual user. For instance, when the vectors of attack are few, cyber attacks are harder to effectively launch, making the cyberspace environment more secure. When the number of attacks is low, the probability that a user will be subject to a cyber attack is less, again making cyberspace more secure. Likewise, when the damage caused by a successful attack is low, the cost of a successful cybercrime for an individual is less severe, meaning the environment is less threatening overall. In every case, as the vectors, occurrence or damage of cyber attacks goes up, the overall security of cyberspace from a user's perspective goes down.

This chapter operationalizes the concept of the vectors of cyber attack via the following measures:

- new vulnerabilities;
- malicious web domains;
- zero-day vulnerabilities;
- new browser vulnerabilities; and
- mobile vulnerabilities.

The concept of the number of attacks is operationalized via:

- botnets; and
- recorded web-based attacks.

Table 2: Summary Statistics for the Security of Cyberspace

	Minimum	Maximum	Mean	Standard Deviation
New Vulnerabilities	4,814	6,787	5,749	781.880
Malicious Web Domains	29,927	74,000	53,317	13,769.99
Zero-day Vulnerabilities	8	24	14.85714	6.336
New Browser Vulnerabilities	232	891	513	240.570
Mobile Vulnerabilities	115	416	217.35	120.85
Botnets	1,900,000	9,437,536	4,485,843	2,724,254
Web-based Attacks	23,680,646	1,432,660,467	907,597,833	702,817,362
Average per Capita Cost	188	214	202.5	8.893818078
Organizational Cost	5,403,644	7,240,000	6,233,941	753,057
Detection and Escalation Costs	264,280	455,304	372,272	83,331
Response Costs	1,294,702	1,738,761	1,511,804	152,502.2526
Lost Business Costs	3,010,000	4,592,214	3,827,732	782,084
Victim Notification Costs	497,758	565,020	523,965	30,342

And the concept of the damage of attacks is operationalized via:

- average cost per data breach;
- overall organizational cost from data breaches;
- the cost of detecting a data breach and escalating;
- post-breach reaction costs;
- lost business costs; and
- victim notification costs.

Table 2 presents some basic summary statistics on the various indicators of the insecurity of cyberspace.

Vectors of Attack

New vulnerabilities are exploitable points in the software code underwriting a program that can provide a cybercriminal with unwanted access to a device.⁷ New vulnerabilities are distinct from zero-day vulnerabilities in that they are publicly known. Companies provide routine updates to their programs (Microsoft updates roughly every Wednesday, for example). These updates often include patches for newly discovered vulnerabilities. Failure to update a program can lead to serious problems, as cybercriminals can exploit peoples' sluggish behaviour to infect a system through these publicly known, but inadequately patched, weak points. Data on new vulnerabilities from 2008 to 2014 are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat*

Reports (Norton Symantec 2009; 2010; 2011; 2012; 2013; 2014; 2015).

Malicious web domains are domains that have known bits of malicious code embedded within them. This code is designed to infect a visiting user's computer with a virus. Malicious web domains are a passive vector of attack for cybercriminals because they require that the user go to an infected domain. Nevertheless, this can still be a potent avenue of attack. Data on malicious web domains are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

New zero-day vulnerabilities are vulnerabilities in software code that are as of yet unknown. The "zero day" part of the name refers to the fact that there have been zero days available to provide a patch that fixes the vulnerability. Zero-day vulnerabilities are fairly rare and quite valuable. Cybercriminals that gain access to a zero-day vulnerability can attack computers easily, as there is no defence against this exploitation; therefore, they are a highly potent vector of attack. Data on zero-day vulnerabilities are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

New browser vulnerabilities are weak points in the code of web browsers, such as Google, Safari and Internet Explorer. As most of the top level of the Internet is digested via a web browser, they are useful avenues for attack by cybercriminals. The data on web browser vulnerabilities

⁷ In the case of the various vulnerabilities discussed in this chapter, the numbers are a count of the new vulnerabilities for that year and not a count of all the vulnerabilities that have ever been discovered.

are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).⁸

New mobile vulnerabilities refer to vulnerabilities that are specific to mobile devices, such as Android devices or iPhones, rather than laptops and desktop computers. The data on mobile vulnerabilities are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

Occurrence of Cyber Attacks

Botnets are computers that have been infected by a virus that allows them to be hijacked and used remotely by a third party for some illicit purpose. Botnets are often employed in distributed denial of service (DDoS) attacks, which require that a large number of requests be made of a website in a short period of time. Botnets are also often used to send spam emails. To become a part of a botnet, an online device needs to have been the subject of a cyber attack. A measure of botnet computers is one way to get at the number of victims of a crime, although certainly not the only one. The number of botnet computers, therefore, gives a sense of the occurrence of successful cyber attacks. Data on botnets are taken from the 2009 through 2015 Norton Symantec *Internet Security Threat Reports* (ibid.).

Recorded web-based attacks are cyber attacks that were launched against one part of the network from an online source and are a good measure of the occurrence of cyber attacks. These attacks exclude cyber attacks that result from, say, the use of an infected USB key. Web-based attacks provide a picture of the overall occurrence of cyber attacks, although, due to reporting problems and the fact that cybercriminals often try to have their attacks go unnoticed, the actual number of attacks is probably higher than the recorded figure. Data on web-based attacks are drawn from the IT security firm Kaspersky Lab's "Security Bulletin" reports (Kaspersky Lab 2008; 2009; 2010; 2011; 2012; 2013; 2014).

The Damage of Cybercrime

The concept of the damage done by cybercrime is operationalized in five ways. This chapter focuses exclusively on the cost of data breaches for companies, although this is certainly not the be-all and end-all of the costs cybercrime imposes onto users of the Internet.

⁸ There is a major discrepancy in the Norton 2013 report compared to the Norton 2014 report. The 2013 report indicates that new browser vulnerabilities rose from 591 in 2011 to 891 in 2012 before falling to 351 in 2013. The 2014 report indicates that new browser vulnerabilities rose from 351 in 2011 to 891 in 2012 before declining to 591 in 2013. The chapter retains the earlier, 2013, data because it actually works against the hypothesis that the security of cyberspace is better than the absolute numbers by moving a higher number earlier in time. In the tests below, using the 2014 data only changes the magnitude, and not the direction, of the relationship.

All the data on breaches is taken from the Ponemon Institute's annual *Cost of Data Breach Study*, which records the overall cost of data breaches (Ponemon Institute 2011; 2013; 2014). Unfortunately, the Ponemon Institute only started collecting a global sample in 2013 and previously only collected the costs associated with US data breaches. The United States is still in the later global assessments, so for the purpose of over-time comparability, only the US numbers are included in the analysis below. Due to the overall lack of statistics on data breach costs, this chapter makes the assumption that the US cost of cybercrime data is indicative of the world's costs. In reality, the average costs for the world are almost certainly far lower than the US costs. For example, in 2013, the organizational cost of data breaches in the United States was \$5,850,000. Globally, the average based on the weighted numbers from the later Ponemon Institute studies, including the United States, is \$2,282,095, or a difference of over twice as much. Using the US numbers, in other words, will overstate the costs of cybercrime and actually work against the argument herein that the security of cyberspace is better than the impression given by the absolute numbers.

Before turning to a discussion of the various measures used to operationalize the cost of cybercrime, it is important to note two additional limitations to the statistics collected on data breaches. The companies studied vary from year to year, as does the number of companies that are observed. Clearly, from a methodological point of view, this is not ideal, as the shifting foundational sands of the studies mean that the inter-year samples are not strictly comparable. Another limitation is that the studies exclude "mega breaches," or those involving more than 100,000 breached records in a single attack. This restriction essentially excludes high-damage but low-probability events in favour of the more representative high-probability but comparatively low-damage events that occur most of the time. Despite all these limitations, the Ponemon Institute's studies of the cost of data breaches are the best publicly available data on the overtime costs of data breaches.

The first operational measure of the cost of cybercrime is the average cost for a company per breached record. This measure shows the organization's cost divided by the number of compromised files. This measure is one way to show how much an organization has to pay as a consequence of cybercrime.

Another way to portray this cost — and the second measure of the costs of cybercrime — is the overall average organizational cost of data breaches in a given year. This figure is basically the total price tag of dealing with data breaches. It is a good measure of the cost of cybercrime because it quantifies the absolute cost that a company needs to pay as a result of online criminal behaviour.

A third measure of the costs of cybercrime involves a company's detection and escalation costs. Data breaches are

bad; undetected data breaches are worse. Companies invest considerable resources into IT security so that they can detect data breaches, and, if warranted, act to repel them, although these sums are not necessarily sufficient. This is a good measure of the cost of cybercrime because it involves the investment that companies need to undertake since they operate in an environment with less than perfect security.

A fourth measure is the cost that an organization needs to pay after a data breach in order to fix any damage done. Cybercrime can often result in damage to software and computer hardware. This is a good measure of the cost of cybercrime, because, like a broken window after a burglar breaks into a person's home, the damage done by cybercrime is not just a result of what is stolen.

A fifth measure of the costs of cybercrime is the cost of lost business. Companies, in particular those that provide an online service, rely on the public's perception that their services are trustworthy. If the public thinks that using a company's services will lead to a loss of personal or financial information, individuals are likely to choose other service providers or cease that activity entirely. The cost of lost business as a result of the occurrence of data breaches is a good measure of the sort of second-order effect of cybercrime on a company's balance sheet.

A final measure of the costs of cybercrime is the cost of notifying victims that their records, be they personal, financial or otherwise, have been compromised in a data breach. Even though companies might have an incentive to cover up a data breach for fear of losing business, many are legally obliged to inform those individuals that have had their information compromised.

TRENDS IN THE VECTORS OF ATTACK: VULNERABILITIES AND MALICIOUS SITES

This section compares the absolute numbers for the various vectors of attack against the normalized trend. In every case, the normalized trend presents a picture of the security of cyberspace that is better than the one presented by the un-normalized absolute figures.

This section looks at vectors of cyber attack, which are basically the ways in which cyber attacks can occur to an Internet user. The relative number of ways in which an Internet user can be attacked are declining, given the growing size of the Internet. One way to think of this is to imagine a city with a number of high-crime neighbourhoods. If the city is made up of 10 neighbourhoods and five of them are dangerous, then the crime rate is 50 percent. If the city grows (as cyberspace has grown) faster than the number of bad neighbourhoods, then the crime rate declines and people are relatively safer. Imagine the hypothetical city grows in size to

15 neighbourhoods, but the number of high-crime areas stays at five. The new crime rate is only 33 percent. The city is safer as a result and a person's chance of being subject to a crime declines. Cybercrime vectors are like the high-crime neighbourhoods.

The analysis below undertakes a number of different normalizations for each measure of the security of cyberspace. A justification for each normalization is provided in each section. Multiple normalizations are used, rather than just a single one for each measure of cybercrime, because there is not an agreed-upon denominator that makes the most sense across the different measures. So, in the interest of painting the broadest possible picture and of forestalling the notion that this chapter uses only the normalizations that support its argument, several normalizations per cybercrime measure are included.

Figure 1 normalizes new vulnerabilities as a vector of attack around the number of Internet users, the number of email users and the number of websites. Since vulnerabilities are weaknesses in computer code, the ideal denominator for new vulnerabilities would be the number of software programs that are in use around the world. Unfortunately, the number of programs is not even partially known. In the absence of this data, Internet users, email users and websites will have to do. The number of Internet users gives an (admittedly partial) impression of the number of devices that are operating online and so indicates the chance that a device will be using software that is afflicted by a new vulnerability. The number of email users is another measure of active devices online, pointing to the odds that a device will be running a flawed program. Finally, websites are hosted using various software programs, all of which can have unexpected vulnerabilities. The number of websites, therefore, provides a measure of the points of interaction online that are operating software that could be prone to cyber attack due to a new vulnerability.

In Figure 1, the trend in the absolute figures suggests that the number of new vulnerabilities is actually worsening between 2008 and 2014, rising from 5,562 new vulnerabilities in 2008 to 6,549 new vulnerabilities in 2014; an increase of 17.75 percentage points over the five years. In contrast, each of the normalized trends suggests that this vector of attack is actually improving over time. For instance, new vulnerabilities normalized around the number of Internet users, a proxy for online devices in this case, fell from 3.56 new vulnerabilities per 1,000,000 Internet users in 2008 to 2.24 vulnerabilities per 1,000,000 Internet users in 2014. This drop amounts to a percentage change of 37.13 percent. In other words, the normalized numbers suggest that the security of cyberspace is greater than what is suggested by the absolute numbers. Indeed, the absolute numbers indicate that the situation is worsening, while the normalized figures actually indicate that the situation is improving.

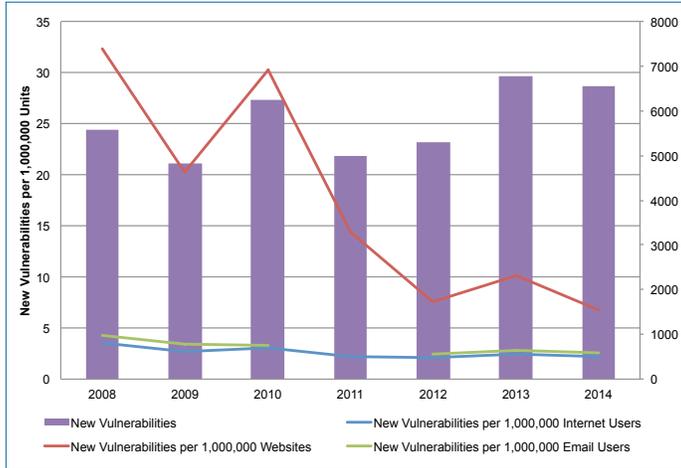
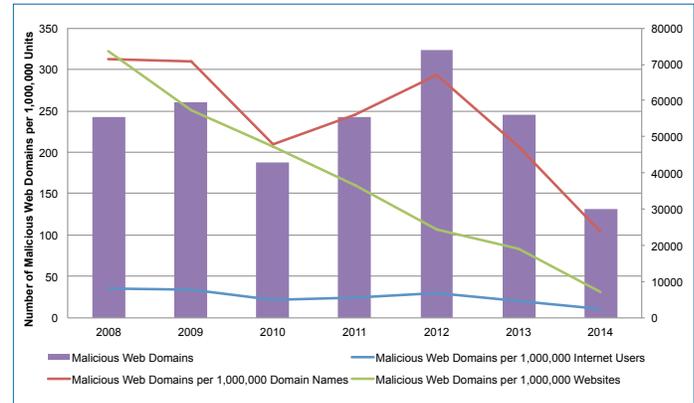
Figure 1: New Vulnerabilities**Figure 2: New Malicious Web Domains**

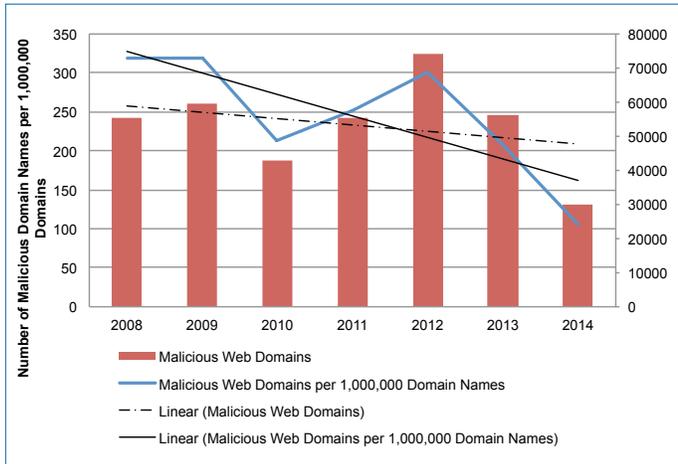
Figure 2 compares the normalized trend among malicious domains as a vector of attack against the absolute number of malicious domains. The number of malicious web domains is normalized around the number of Internet users, the number of web domains and the number of websites. Clearly, the most natural data manipulation is to normalize malicious domains around the total number of domains (which is done in both Figure 2 and then in more detail in Figure 3). Normalizing malicious domains around the number of Internet users makes sense because the latter measures the number of people that can be affected by a malicious domain, which shows the trend in potential infection rates. As mentioned above, the number of web domains is a smaller subset of the total number of websites, which can have subsidiary pages and the like. Normalizing the number of malicious web domains around the number of websites provides another glimpse of how problematic a given number of malicious web domains are likely to be because it shows how many websites might be affected and so how many webpages might be a threat to the security of cyberspace.

As shown in Figure 2, the number of absolute new malicious web domains has remained fairly constant over time, with an initial increase from 2010 to 2012 being followed by a decline from 2012 to 2014. In contrast to these fairly stable numbers, the normalized trends in malicious web domains per 1,000,000 Internet users and per 1,000,000 websites both strongly point toward an improving security situation in cyberspace. However, probably the most appropriate normalization in this case is the number of malicious web domains per 1,000,000 Internet domains, since the basic unit of measure (domains) is the same. Here, the absolute number of malicious domains and the normalized trend track together fairly consistently, but the actual trend underlying the two sets of data shows a clear difference in degree.

Figure 3 looks at just the comparison of the absolute number of malicious domain names and the trend in malicious domains normalized around the total number of domains. The appearance that these two indicators track together over time suggests that there is a fairly static proportion of all web domains that are malicious. However, this initial impression is misleading in the sense that the two sets of numbers are changing at very different speeds. The two trend lines in Figure 3 show that between 2008 and 2014 both the absolute and the normalized trends have been improving. Comparing the rate at which the situation is improving tells a different story. The absolute number of new malicious domains has fallen from 55,389 malicious domains in 2008 to 29,927 malicious domains in 2014, a decline of 45.96 percent. In contrast, the normalized numbers fell from 312.93 malicious domains per 1,000,000 domains in 2008 to only 103.91 malicious domains per 1,000,000 domains in 2014, which amounts to a decline of 66.79 percentage points. As with the new vulnerabilities, the data from Figures 2 and 3 support the idea that the absolute numbers overrepresent the insecurity of cyberspace compared to the normalized trends by showing the picture improving more slowly than is actually the case.

Figure 4 presents the data on the number of zero-day vulnerabilities normalized around the number of Internet users, web domains and the number of websites, and contrasts these numbers with the absolute trend. As with new vulnerabilities, the best measure to normalize zero-day vulnerabilities would be the number of software programs used in the world, the data for which does not exist. Nevertheless, since zero-day vulnerabilities are weaknesses in computer code, the normalization that makes the most sense is the number of zero-days per 1,000,000 websites, since websites rely on a growing number of software platforms (think of the Heartbleed zero-day exploit in Secure Sockets Layer [SSL] in 2014). In the interest of presenting the broadest possible story, the number of zero-day vulnerabilities normalized around the number of Internet users and email users are also included

Figure 3: Normalized versus Absolute Domains



(both proxies for the number of potentially vulnerable devices operating various pieces of software).

The dotted trend line in Figure 4 shows that over time the absolute number of zero-day vulnerabilities is getting larger, suggesting a worsening cyber security environment. This finding is mirrored by the trend in zero-day vulnerabilities per 1,000,000 email users and per 1,000,000 Internet users. However, the trend in zero-day vulnerabilities per 1,000,000 websites is actually declining over time, despite a jump upward in 2013. To the extent that normalizing the number of zero-day vulnerabilities around the number of online websites is the most accurate measure of this vector of cyber attack, the fact that the trend is negative suggests that, as is the case with the other measures, the security of cyberspace is improving over time even as the absolute number of zero-day exploits increases.

Figure 5 summarizes the data on browser vulnerabilities as a vector of cyber attack, depicting both the absolute numbers and the number of new browser vulnerabilities normalized around the number of Internet users, the number of websites and the number of Google searches. The number of new browser vulnerabilities are normalized around the number of Internet users because this manipulation of the data shows the rate at which people will come into contact with vulnerable browsers (not accounting for the fact that different browsers are used more frequently than others). The number of new browser vulnerabilities is normalized around the number of websites because these are the points of online interaction that people are trying to reach via a web browser. The more websites that exist, the more people will be pulled to use a web browser and so the larger the potential that a browser will affect an online device. Finally, in what is probably the most accurate normalization, the number of browser vulnerabilities is divided by the number of Google searches. Google searches capture the frequency

Figure 4: New Zero-day Vulnerabilities

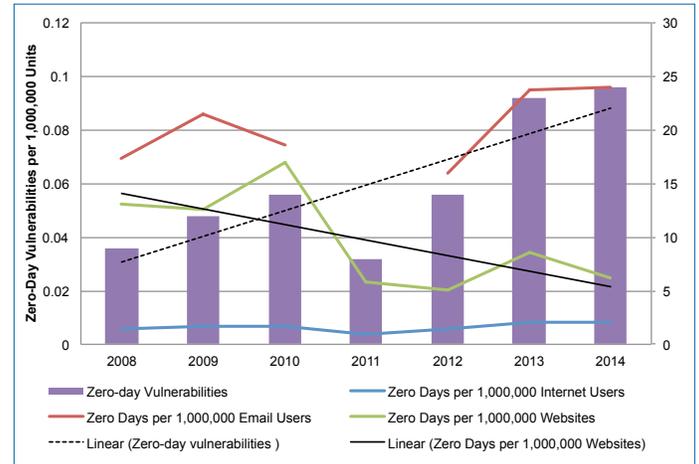
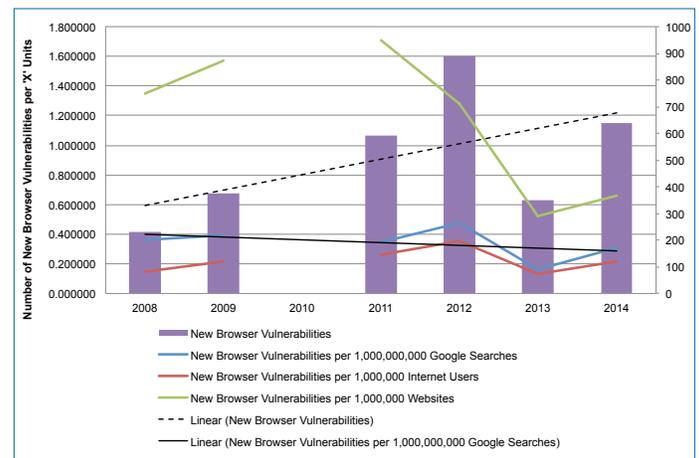


Figure 5: New Browser Vulnerabilities



with which a globally dominant web browser is actually being used and thus how probable it is that an Internet user will come into contact with a vulnerable browser.

As shown by the dotted trend line in Figure 5, the absolute number of new browser vulnerabilities is generally increasing over time, with 639 browser vulnerabilities in 2014 compared to 232 in 2008 (an increase of 175 percentage points). New browser vulnerabilities normalized around the number of Internet users is also slightly escalatory over the full seven-year period. In contrast, new browser vulnerabilities as a proportion of all websites show a generally de-escalatory trend and an improving cyber security situation. Most telling, given its likely accuracy as a measure of effect of new browser vulnerabilities, the number of vulnerabilities normalized around Google searches is negative, as shown by the solid black trend line. In numerical terms, the number of new browser vulnerabilities per 1,000,000,000 Google searches drops from 0.364 new vulnerabilities per 1,000,000,000 Google searches in 2008 to 0.305 new vulnerabilities per 1,000,000,000 Google searches in 2014, a decline of

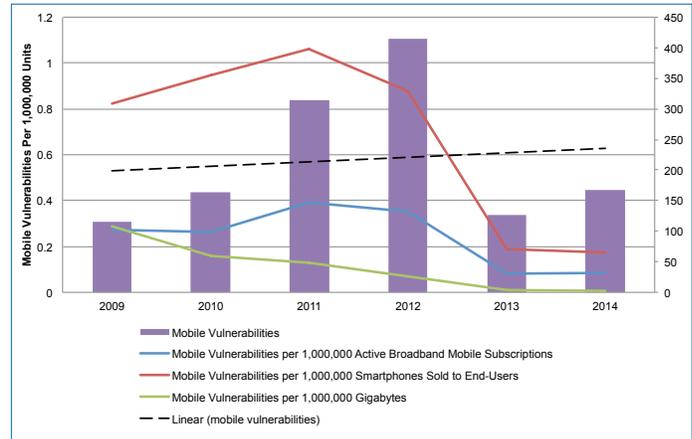
16.23 percentage points. Overall, the numbers on new browser vulnerabilities as a vector for cyber attack again support the idea that the absolute numbers paint a worse picture of the security of cyberspace than the normalized numbers. In this case, the absolute numbers indicate that the situation is worsening, while the normalized numbers say that things are actually improving.

Finally, Figure 6 shows the number of new mobile vulnerabilities and the number of new mobile vulnerabilities normalized around the number of active broadband mobile subscribers, the number of smartphones sold to end-users, and the volume of mobile data usage in gigabytes. These three normalizations make eminent sense because mobile vulnerabilities (glitches and weaknesses in the operating system or associated software of mobile devices) can only affect mobile users. Each normalization helps clarify the real risk that a user faces when using a mobile device to access the Internet. Normalizing new vulnerabilities around active mobile broadband subscriptions shows how likely a user is to be affected by a new vulnerability. Normalizing the number of new vulnerabilities around the number of smartphones sold to end-users shows the likelihood that a particular device will be afflicted by a cybercrime. Finally, normalizing the number of new mobile vulnerabilities around the volume of mobile traffic shows how problematic weaknesses are in light of how much people use mobile platforms to access the Internet.⁹

As shown in Figure 6, mobile vulnerabilities have expanded rapidly since 2009, with the number of new mobile vulnerabilities increasing from 115 in that year to 415 at the peak in 2012, before declining to 127 in 2013 and jumping up again to 168 in 2014. This growth in mobile vulnerabilities tracks the growth in the use of mobile devices, both in the developed world and among new entrants to the Internet. From 2009 to the peak (in terms of new mobile vulnerabilities) in 2012, the absolute numbers indicate that the number of new vulnerabilities rose by 261 percentage points. Across the whole sample, the absolute numbers on new mobile vulnerabilities indicate that the security of cyberspace is growing worse over time, even with the significant drop in new vulnerabilities in 2013, as shown by the long-dashed trend line. In contrast, the three normalized measures each show that the security of cyberspace is actually improving. The reduction in new vulnerabilities relative to the various measures is also substantively large. For example, the number of new vulnerabilities per 1,000,000 gigabytes of mobile data fell from 0.29 vulnerabilities per 1,000,000 gigabytes in 2009 to 0.0064 vulnerabilities per 1,000,000 gigabytes in 2014, a reduction of roughly 97.7 percentage points. Active

⁹ Clearly, the best measure in this case would be if both vulnerabilities and broadband subscriptions specified the type of operating system or software that was problematic and used on the device. Since this data does not exist, the data included in the text is the next best option.

Figure 6: New Mobile Vulnerabilities



mobile broadband subscriptions, for their part, fell from 0.273 new vulnerabilities per 1,000,000 subscriptions in 2009 to 0.086 vulnerabilities per 1,000,000 subscriptions in 2014, a reduction of 68.43 percentage points. Finally, the number of new vulnerabilities per 1,000,000 smartphones sold fell from 0.826 in 2009 to 0.173 in 2013, a reduction of 79.02 percentage points. Clearly, the normalized numbers paint a radically different picture of the security of cyberspace than the absolute numbers, the latter showing the situation getting worse and the normalized numbers showing the situation rapidly improving. In short, mobile vulnerabilities continue to grow, but they are growing more slowly than the actual use of mobile devices. Essentially, the absolute numbers say that the situation is worsening, when, as shown by the normalized numbers, the security of cyberspace is actually improving.

When it comes to the potential vectors of cyber attack, the security of cyberspace is far better than what is shown by just looking at the absolute numbers. In four of the five vectors of attack (new vulnerabilities; zero-day exploits; browser vulnerabilities; and mobile vulnerabilities), the absolute numbers say that the situation is getting worse over time, while the normalized numbers show the opposite: cyberspace is becoming more secure. In the remaining case (malicious domains), both the absolute and the normalized numbers indicate an improving situation, but the former shows cyberspace getting better at a slower rate than the latter. In short, when it comes to vectors of attack, cyberspace is a lot safer than one might think.

OCCURRENCE OF CYBER ATTACKS: WEB-BASED ATTACKS

This section looks at the occurrence of cyber attacks in absolute terms compared to the normalized trend in the number of botnet computers and cyber attacks between 2008 and 2014, given the growing size of cyberspace. On botnets, or computers that have been successfully targeted

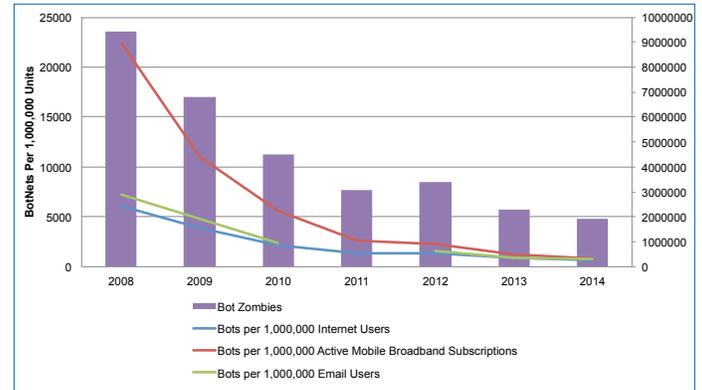
by a cyber attack, both the absolute and the relative numbers show that things are improving over time. The normalized numbers, however, point to a situation that is getting better faster, when compared to the absolute numbers. Both the absolute and the normalized numbers for the occurrence of cyber attacks indicate that the situation has worsened overall since 2008-2009. At the same time, both sets of numbers show the situation improving since 2013 (in the case of the absolute numbers) and 2012 (in the case of the normalized numbers). Yet, the normalized numbers not only show the situation getting better sooner, but also indicate that things are getting better faster, when the growing size of cyberspace is taken into account. Looking at the actual occurrence of cyber attacks, in other words, the absolute numbers again paint a worse picture of the trends than the relative ones.

The occurrence of cyber attacks is like the occurrence of robbery or violent crime in the real world. Cyber attacks directly target Internet users in some way or another, as crime does in the physical space. To be rather selfish about it, you might not really care how much violent crime there is in a city, only your chances of being the subject of that crime. The basic story in cyberspace is that there has been an increase in violent crime in our hypothetical city of 100,000 people since 2008. But, since the early 2010s, the situation has stabilized and even started to improve overall. More pointedly, a person's chances of being the subject of a cybercrime have declined as the size of cyberspaces has grown and the number of attacks has fallen. Things are getting better, even if the golden age of low crime levels seems to have passed.

Figure 7 plots out the absolute number of botnets compared to the number of botnets normalized around the number of Internet users, active mobile broadband subscriptions and email users. These three measures of the size of the Internet mesh well with the nature of botnets. Botnets are hijacked computers, which today can be desktops, laptops, phones, fridges or any other connected device. Once commandeered, these devices can be used to send spam and launch DDoS attacks. To become part of a botnet, a computer needs to become infected with a malicious program. This means that the computer needs to be operational (Internet users, active mobile broadband subscriptions and email users express the number of operational computers, although the number in each case is smaller than the actual number of online devices) and need to be infected somehow (Evans 2011).¹⁰ As such, the three normalizations that make the most sense are botnets divided by online users.

¹⁰ This conceptualization focuses on the risk of having a computer become a botnet and not the other side of the issue of whether a botnet will be used to launch a DDoS attack on a website. Looking from this angle, the normalization of botnets around the number of Internet, active mobile broadband subscriptions or email users expresses how large the criminal element is as a proportion of all users.

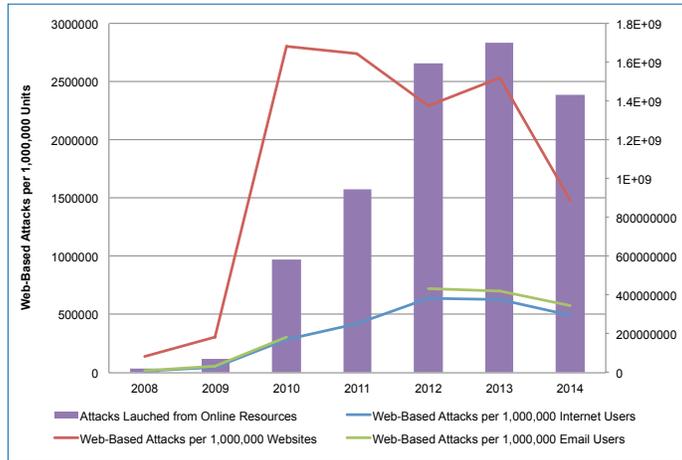
Figure 7: Botnets



As is clear from Figure 7, while both the normalized and the absolute numbers point to a decline in the number of botnet computers between 2008 and 2014, the normalized numbers show a far steeper drop.¹¹ The absolute number of botnet zombies, which is a count of the number of infected computers worldwide, fell from 9,437,536 in 2008 to only 1,900,000 in 2014, which is a drop of 79.9 percentage points. In contrast, the number of botnets normalized around the number of Internet users fell from 6,041.69 botnets per 1,000,000 Internet users to 650 botnet computers per 1,000,000 users during this same period, amounting to a decrease of 89.24 percent. Similar magnitude declines are found for both active mobile subscriptions (-96.3) and email users (-89.5). This data suggests that the absolute figures overrepresent the insecurity of cyberspace compared to the normalized numbers by exaggerating the problem of botnets as a potential vector of cybercrime.

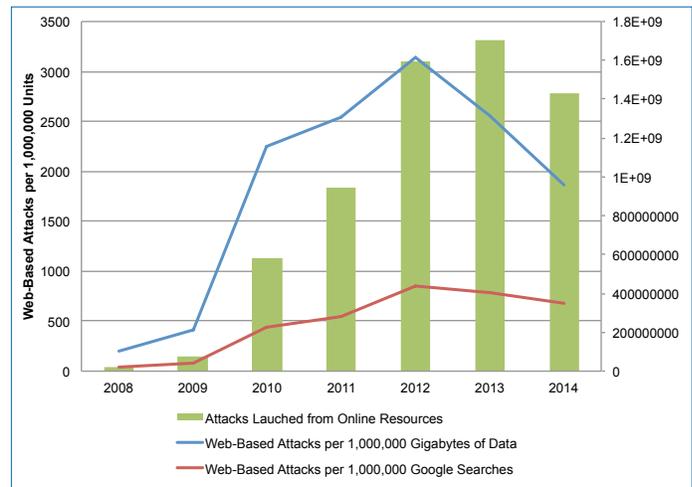
Figure 8 shows the level of absolute web-based attacks compared to the number of such attacks normalized around the number of Internet users, the number of websites and the number of email users. The normalization of the occurrence of attacks around both Internet users and the number of email users captures the idea that cyber attacks target individuals who use the network and that one's chance of being affected by a cybercrime is determined by both the number of attacks and the number of other Internet users. These normalizations, in other words, are similar to normalizing crime statistics around the number of people that live in an affected area. Websites are one clear source of web-based attacks. The normalization of the number of attacks around the number of websites (crudely) shows how frequently attacks occur given the available stock of online points of interaction.

¹¹ The processes for identifying and counting botnets have also improved over time, rendering a more accurate picture of the total number of active botnet computers. While it is impossible to know for sure, it is plausible that earlier counts under-represented the number of botnets, which suggests that the decline has been even steeper. I am grateful to Laura DeNardis for pointing this out to me.

Figure 8: Web-based Attacks

As shown in Figure 8, the absolute numbers point to a strong escalatory trend in cyber attacks, indicating a worse level of security in cyberspace between 2008 and 2014. For example, there were 23,680,646 web-based attacks in 2008 and some 1,432,660,467 attacks in 2014, which is a 5,950 percentage point increase over just seven years. In contrast, the number of web-based attacks per 1,000,000 Internet users has only increased from 15,159.8 in 2008 to 489,756.7 in 2014, which is an increase of only (using that term very loosely) 3,130.63 percent. The normalized trends also all suggest that, while the cyberspace security situation is definitely worse than in 2008 and 2009, the trend in normalized cyber attacks has improved since 2010 in the case of attacks per 1,000,000 websites, and since 2012 in the case of attacks per 1,000,000 Internet and 1,000,000 email users. The absolute numbers suggest that, at best, the situation started to improve only in 2014, although it is possible that the low number of web-based attacks in 2014 is a statistical fluke rather than the start of a real trend in the absolute numbers.

Figure 9 normalizes the number of cyber attacks around the volume of Internet traffic and the number of Google searches for the 2008–2014 period. The intuition behind both normalizations is that, even if there is a constant rate of web-based attacks, the absolute number of attacks should grow as the Internet is used more and more in our daily lives. In such a case, more web-based attacks might not mean an individual user is more likely to be subjected to a cybercrime. What matters is the rate at which web-based attacks occur. Normalizing web-based attacks around the total volume of Internet activity roughly indicates what proportion of Internet activity is actually malicious and aimed at undermining the security of cyberspace. As a caveat, the rapid growth in video streaming likely biases these numbers, as streaming video takes up a lot of bandwidth and does not usually come with the same

Figure 9: Web-based Attacks and Internet Traffic Flows

level of security risk as generic web surfing.¹² Normalizing the occurrence of web-based attacks around the number of Google searches is another way to get at the rate at which online activity is likely to be marred by cybercrime. In this case, the measure of online activity is imperfect because Google searches are only a significant subset of all search engine queries and do not encompass all online activity.¹³

As shown in Figure 9, both the absolute numbers and the normalized trends point to an overall escalatory situation in the occurrence of cyber attacks between 2008 and the end of 2014. Yet, there is some hope as web-based attacks fell from 1,700,870,654 attacks in 2013 to 1,432,660,467 attacks in 2014. This amounts to a decline of around 15.77 percent. In contrast, these data show that the normalized trends both start to improve sooner (2012 rather than 2013) and fall more sharply than the absolute numbers. The number of web-based attacks as a share of all Internet traffic, for example, falls from roughly 3,143 attacks per 1,000,000 gigabytes of data in 2012 to roughly 1,868 attacks per 1,000,000 gigabytes of data in 2014, which amounts to a decline of 40.55 percent. The number of web-based attacks normalized around the number of Google searches likewise falls from roughly 852 attacks per 1,000,000 Google searches in 2012 to 684 attacks per 1,000,000 Google searches in 2014, or a decline of 19.7 percentage points. In short, looking at attacks as a proportion of data flow and online activity, the security of cyberspace is again improving both sooner and faster than what is shown by the absolute numbers.

¹² I am grateful to the reviewer for pointing out this limitation in the data.

¹³ A better measure that is not publicly available would be web queries, where people are making requests to view websites. Again, I am grateful to the reviewer for pointing out this potential measure. I only lament that I could not find the data to bring the idea to fruition.

There has indeed been a massive increase in the absolute number of web-based cyber attacks since 2008. Yet, while the glory days of 2008 and 2009 might be gone, since 2010–2012, the rate at which web-based cyber attacks have occurred has declined a lot more than one might otherwise think when factoring in the growing size of the Internet. All five normalized trends bear out this claim.

Overall, the findings in this section show that, when compared to the absolute numbers, the various normalized numbers all point to a situation that both starts improving sooner and that improves more rapidly. The security of cyberspace, in other words, is better than one might think looking at just the absolute numbers.

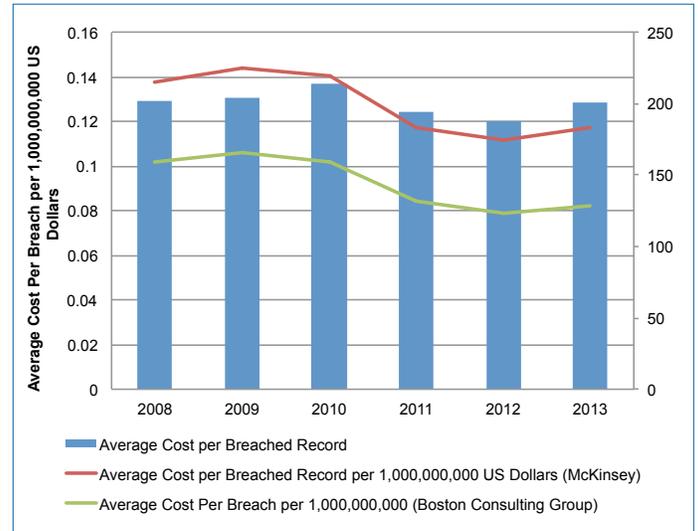
THE COST OF SUCCESSFUL CYBER ATTACKS

This section compares the absolute numbers to do with the various costs of data breaches with the same numbers normalized around the size of the Internet's contribution to the global economy. Underlying this move is the idea that we need to understand the cost of cybercrime relative to the economic benefits that accrue from the Internet. The real concern would be when the costs of doing business are greater than the benefits produced by using the Internet as a platform for communications and commerce, as firms would then opt out of the system. Normalizing the numbers in this way shifts the question from what a firm pays as a result of data breaches to what sort of economic damage is done in general terms by cybercrime compared to the benefits that are generated by the Internet economy. Again, the absolute numbers consistently suggest a worse cyberspace environment than the normalized numbers.

When it comes to the costs of cybercrime, the value added of the Internet is outpacing the costs that Internet-enabled cybercrime imposes on society. In other words, in net terms, having the Internet is still beneficial, even though cybercrime inflicts economic damage. In the daily world, another example of a sort of dual-use system that both generates economic growth and facilitates crime is the global financial system, which can be used to provide loans and transfer funds, but which can also be used to launder money and avoid taxes. At a social level, what matters are net gains, and, in the case of the Internet and cybercrime — as in the case of the global financial system — things are looking pretty good.

Figure 10 looks at the average cost per breached record in absolute terms compared to these numbers normalized around both the McKinsey & Company and the Boston Consulting Group's estimates for the Internet's contribution to global GDP. The absolute numbers paint an image of a roughly constant average cost per breached record, with the

Figure 10: Average Cost per Breach



cost in 2008 being \$202.00 and \$201.00 in 2013.¹⁴ In contrast, both sets of normalized figures show a reduction over this same time period. The numbers normalized around the McKinsey & Company estimates for how much the Internet contributes to the global economy show a drop in the average cost per breached record from \$0.14 cents per \$1,000,000,000 dollars of Internet contribution in 2008 to \$0.12 cents per 1,000,000,000 in 2013. This decline amounts to a 14.85 percentage change in the normalized cost per breached record. Likewise, the numbers normalized around the Boston Consulting Group estimates show a similar declining trend, with the average cost per breached record per \$1,000,000,000 of the Internet's contribution falling from \$0.10 cents in 2008 to \$0.08 cents in 2010 (a reduction of 19 percentage points). The comparison of the data on the average cost per breached record indicates that the absolute trend depicts a relatively constant level of cost, while the normalized trends show a decreasing cost. Overall, the absolute figures overrepresent the cost of cybercrime in this area compared to the normalized figures.

Figure 11 presents data on the average overall organizational cost that a company is forced to bear as a result of data breaches. A pretty consistent message emerges across all the numbers, with the absolute and normalized trends pointing to a declining cost due to data breaches and thus an overall improvement in the security of cyberspace. However, a comparison of the rate at which the numbers are declining paints a slightly different picture. For the absolute figures, the overall organizational cost fell from a high of \$7,240,000 in 2010 to just \$5,850,000 in 2013. This drop amounts to a decrease of 19 percentage points. In contrast, looking at the Boston Consulting Group's

¹⁴ The addition of a trend line shows a slight decline in the absolute numbers over the full sample.

Figure 11: Organizational Cost

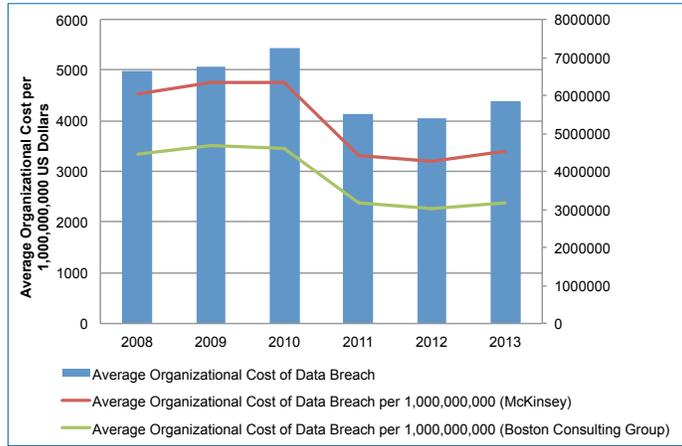
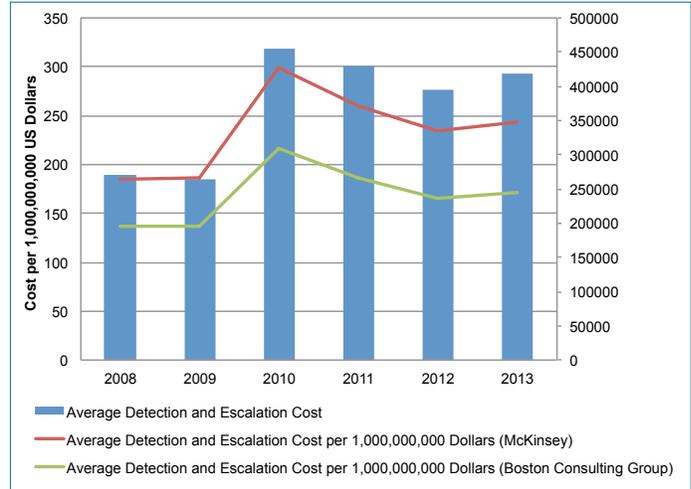


Figure 12: Average Detection and Escalation Costs



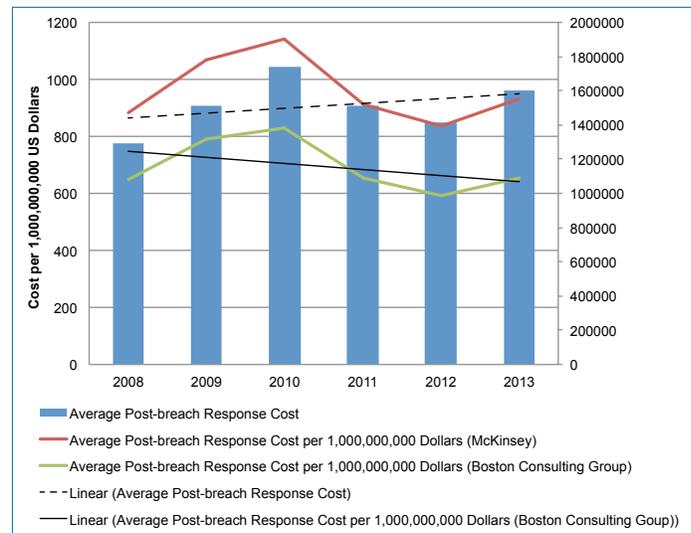
estimates for the size of the Internet’s contribution to GDP, the number falls from a peak value of \$3,513.60 for every billion that the Internet contributed to global GDP in 2009 to a low of \$2,390.19 per \$1,000,000,000 in 2013, amounting to a drop of 32 percentage points.¹⁵ In short, organizational costs due to data breaches are declining across both data forms, but the rate of that decline varies.

Overall, the comparison of the absolute and normalized cost of detection and escalation shows that, since 2008, the costs have uniformly increased, but that the absolute numbers have registered a larger percentage increase in that time compared to the normalized numbers. Likewise, since the high point in terms of the costs in 2010, the absolute numbers show a smaller decline in the costs of detection and escalation compared to the normalized trends. Once again, the absolute numbers paint a more dismal picture of the costs of cybercrime than the normalized figures, suggesting that the security of cyberspace is actually greater than is commonly perceived.

Figure 12 compares the absolute and normalized costs associated with detection and escalation in response to a data breach. In this case, all three sets of numbers point to growing detection and escalation costs since 2008, with a slight reduction in the costs since 2010. Once again, focusing on the magnitude of the changes provides interesting nuance to the picture. From 2008 to 2013, the absolute cost of detection and escalation rose from \$271,084 to \$417,700 or an increase of 54.1 percentage points. In contrast, the numbers normalized around the Boston Consulting Group estimates of the size of the Internet’s contribution to global GDP show that the costs have increased from \$136.17 per \$1,000,000,000 in 2008 to only \$170.66 per \$1,000,000,000 in 2013. This change amounts to only a 25 percentage point increase over that time period. In short, the normalized trends show that the growth in the costs of escalation and detection is less pronounced compared to the absolute figures. A similar story of different magnitude changes emerges if we look at the drop from the high point of detection and escalation costs in 2010 compared to the costs in 2013. Here, the absolute values decline from \$455,304 in 2010 to \$417,700 in 2013, or a decrease of roughly 8.3 percentage points. In contrast, the numbers normalized around the Boston Consulting Group estimates decrease from \$216.93 per \$1,000,000,000 in 2010 to \$170.66 per \$1,000,000,000 in 2013, which amounts to a reduction of roughly 21 percentage points.

Figure 13 presents data on the absolute and normalized trends in post-breach response costs. At first blush, both the absolute and the normalized numbers paint a roughly consistent picture. A more in-depth comparison reveals two points that suggest the absolute numbers overrepresent the post-breach response costs of cybercrime. First, the absolute

Figure 13: Post-breach Response Costs



15 The McKinsey & Company numbers also suggest a larger decline for the normalized trend of around 28.5 percentage points.

numbers indicate an escalatory trend in costs, as shown by the long-dashed trend line. In contrast, the numbers normalized around the Boston Consulting Group estimates for the Internet’s contribution to the global economy show a de-escalatory or declining trend, as shown by the solid black trend line (the McKinsey & Company numbers also point to a declining trend). Secondly, the rate at which the post-breach costs have declined since the high-water mark of 2010 into 2013 shows a greater decline for the normalized numbers compared to the absolute numbers. In particular, the absolute costs fell from \$1,738,761 in 2010 to \$1,599,996 in 2013 or a decrease of 7.98 percentage points. In comparison, the numbers normalized around the Boston Consulting Group’s estimates show a decline from \$828.44 per \$1,000,000,000 in 2010 to \$653.73 per \$1,000,000,000 in 2013, which amounts to a decrease of 21.1 percent. With respect to the post-breach response costs, the absolute numbers point to both a worsening situation and a slower rate of potential improvement, while the normalized numbers point toward a generally improving situation and a larger decrease since the highest level of costs in the sample.

Figure 14 looks at the costs that firms need to endure due to lost business after they have been subject to a data breach. All three sets of numbers show a declining trend in terms of the lost business costs, which could suggest consumers are getting used to data breaches as a part of business in the digital age or that businesses are becoming more adept at managing the public relations side of data breaches. Running a comparison of the rate at which the costs have declined shows again that the absolute numbers depict a comparatively worse environment compared to the normalized trends. For example, in absolute terms, the lost business cost due to cyber attacks faced by firms in 2008 was \$4,592,214. By 2013, that number had declined to \$3,324,959. The percentage change in the absolute numbers amounts to a decrease of 27.6 percentage points. The numbers normalized around the Boston Consulting Group’s estimates for the Internet economy fell from \$2,306.74 per billion in 2008 to \$1,358.51 per billion in 2013. This change amounts to a decrease of 41.1 percentage points. Once again, the normalized numbers point to a situation where the lost business costs suffered by firms are improving faster than the costs as they are suggested by the absolute numbers.

Figure 15, finally, presents data on the normalized and absolute trends in the costs that companies need to incur to inform individuals that their data has been breached. Here, despite the significant drop in the absolute cost of notification from \$565,020 in 2012 to \$509,237 in 2013, the general trend in the absolute numbers is toward higher and higher notification costs, as evidenced by the long-dash trend line in Figure 15. In contrast, the trend in both the normalized figures suggests that notification costs are actually declining between 2008 and 2013. In this case, the

Figure 14: Lost Business Costs

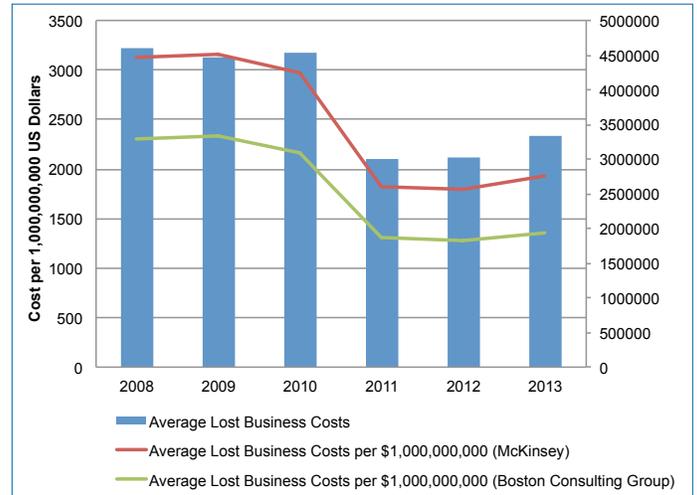
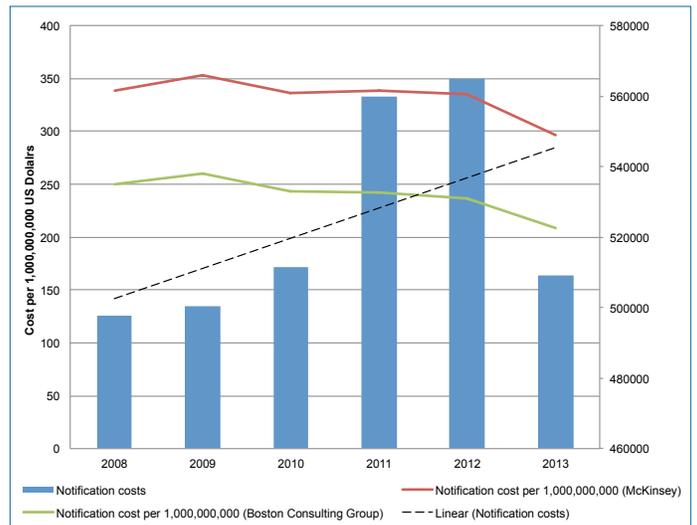


Figure 15: Notification Costs



absolute numbers paint a picture of an increasingly costly security environment, while the normalized numbers suggest that the situation is actually getting better.

So, what conclusions can be drawn from these data on the cost of data breaches as a measure of the costs of cybercrime? Basically, the absolute numbers depict a worse cyber security situation than the normalized numbers. As with the measures for the vectors of cyber attack and the occurrence of cyber attacks, the absolute numbers create the perception that the security of cyberspace is worse

than what is actually suggested by the more accurate normalized numbers.¹⁶

A few qualifiers are needed to temper these conclusions. The numbers in these cases are imperfect, as outlined above. Two points are worth reiterating. First, the economic contribution of the Internet to global GDP is likely larger than what is included in this study due to the assumption that the static, one-year estimates found in the McKinsey & Company and Boston Consulting Group studies are constant forward and backward throughout time. Secondly, the cost of data breaches is likely lower than what is found in these data, since the costs of cybercrime in the United States are, at least according to the Ponemon Institute's studies, consistently higher than the global average. Both of these qualifiers would actually strengthen the argument of this chapter by lowering the various costs of cybercrime, while increasing the Internet's contribution to global GDP. Normalizing these lower numbers around this larger contribution suggests that the normalized trends would be even lower still.

In conclusion, in two of the six tests conducted in this section (post-breach response costs and notification costs), the absolute numbers point to a worsening situation, while the normalized numbers actually indicate that costs are declining. In three of the six cases (average cost per capita, overall organizational costs and lost business costs), both sets of numbers point to an improving situation, but the normalized numbers show the situation improving faster than the absolute numbers. Finally, in the last case (detection and escalation costs), both sets of numbers say the situation is getting worse, but the absolute numbers say that things are falling apart faster than the normalized numbers. Taken together, these findings once again indicate that the security of cyberspace (this time in terms of the costs of cybercrime) is actually better than the impression given by the commonly touted absolute numbers.

CONCLUSIONS AND POLICY RECOMMENDATIONS

What are the actual trends in cybercrime? Is the situation getting worse, better or staying roughly the same over time? We currently have a flawed picture of the security of cyberspace. Instead, a more accurate picture requires that the numbers on the occurrence of cybercrime be normalized around indicators that capture the growth and growing importance of cyberspace. To test this proposition, data on various indicators of the size of the cyberspace

were collected, with a particular focus on users, points of interaction and the volume of online activity. Various measures of the occurrence of cybercrime were examined, with a focus on vectors of attack, the occurrence of attack and the cost of attacks. In every instance, the normalized numbers suggest that the security of cyberspace is better than what is found when one looks only at the absolute numbers. If you take lessons from the 13 normalizations, you find that six (almost half) point to a situation where the absolute numbers show a deteriorating situation while the normalized numbers actually show that things are getting better. In another six of the tests, both numbers show the situation as improving, but the normalized numbers usually indicate that things are getting better sooner and faster compared to the absolute numbers. Finally, in the one case where both sets of numbers show the situation worsening, the absolute numbers still indicate that things are getting worse faster than the normalized numbers. Cyberspace, in other words, is more secure than is commonly perceived.

Research Conclusions

Any conclusions drawn from this research need to be qualified in light of the relatively poor data that is available for study. As pointed out above, an irony of cyber security research is that we live in an age of big data, but very little of this data on cyber security trends is actually publicly available. If the data underlying the study is inaccurate or subject to changes, then the conclusions themselves are also in need of revision. One likely scenario is that many of the indicators for cybercrime are probably higher than the data herein indicates. Software vulnerabilities go undisclosed. Cyber attacks go undetected. Data breaches go unreported. Nevertheless, this chapter maintains that cybercrime in its three modalities (vectors, occurrence and costs) needs to be normalized in order to be properly understood, as has been done here. The numbers might be skewed, but they are definitely more accurate than the simple absolute figures.

Some interesting stories emerge when one looks more closely at some of the trends in the various figures. Obviously, the small number of data points restricts the confidence that we can have in any observations, but there are some suggestive tendencies. For instance, the data on botnets in Figure 7 shows that there has been a steady reduction in the number of botnets since 2008, both in absolute terms and as a proportion of the number of Internet users, email users and websites. This decline potentially suggests that people have become more conscious of the danger of having their computer commandeered for nefarious purposes and have taken steps (such as the use of anti-virus software or being more careful about sites visited) to prevent its occurrence. It could also suggest that there has been a more concerted and coordinated international effort by law enforcement

¹⁶ In the future, the absolute average cost of a data breach might steadily increase as more and more companies and state bureaucracies digitize their information. From a corporate or bureaucratic perspective, digitization promises many cost-saving and efficiency advantages. However, it also creates a larger potential cost if a data breach does occur. The future, in other words, might not be well predicted by the current trend of an improving cost scenario.

agencies and private companies, such as Microsoft, to take down existing botnet networks and operators (Europol 2015). The cause of the decline is likely a mixture of both. Law enforcement efforts are knocking botnets offline, reducing the stock of infected computers, and individual actions may be slowing the rate of infection, reducing the growth of new botnets over time.

The absolute and normalized data in Figures 8 and 9 potentially tell an interesting story regarding whether cybercriminals or cyber security providers hold the initiative.¹⁷ From 2009 to 2012, there is a rapid growth in both the absolute and the normalized number of web-based attacks, suggesting that cybercriminals are among the first to recognize the ways in which new technology can be exploited to make a profit. During 2012, the trend starts to reverse itself, and, in 2013 and 2014, both sets of numbers start to decline. This finding suggests two things. First, the Internet is growing rapidly and at a faster pace each year, which explains the rapid drop in the normalized number of attacks. Second, the decline in the absolute numbers also suggests that law enforcement efforts and individually undertaken security measures are effective at curbing the occurrence of web-based attacks. One interesting supposition that follows from this conclusion is that there are likely to be waves of web-based attacks in the future. Cybercriminals might quickly learn how to exploit new technologies, increasing crime, only to be followed by counteraction by individuals, businesses and law enforcement, which results in a decline in web-based assaults. This cyclical pattern, seen in a preliminary way in the data contained here, will likely be borne out as time goes on.

Lastly, as shown in Figures 10 and 11, both the average cost per breached record and the overall organizational cost of data breaches are declining in both absolute terms and normalized terms. Together, these two trends suggest that the number of data breaches overall might be declining, since both the average cost and the overall organizational cost are declining.¹⁸ One limitation to what can be said on the basis of this data is that the available numbers exclude mega breaches, which compromise over 100,000 records in a single attack. It is also possible, therefore, that the costs of low-grade data breaches are declining because the size of your average data breach is increasing. At the same time, the available evidence suggests that most data breaches tend to be small and targeted at small-to-medium size enterprises (SMEs) (Gow, n.d.). In any event, based on the evidence presented here, the cost of data breaches seems to be decreasing.

¹⁷ I am grateful to the reviewer for pointing out this interpretation to me.

¹⁸ I am again grateful to the reviewer for highlighting this interpretation of the data to me.

Overall, these research results suggest that the security of cyberspace is actually better than what people might think from looking just at the absolute numbers. Assessing the precise effectiveness of cyber security measures given these trends is difficult because it requires a clear account of the counterfactual — that is, what would have happened in the absence of such policies. Put another way, an increasing trend might have actually increased even more or a declining trend might have been less pronounced had a particular policy not been in place. Despite this limitation, one conclusion that can be drawn from the presented evidence is that current cyber security efforts are effective enough to limit the growth in vectors of attack, occurrence of attacks and the costs of attack to some extent. Since these signs of insecurity in cyberspace are not worsening too quickly in most cases, the rapidly growing size of cyberspace actually means that the overall security of cyberspace is, in a lot of cases, generally improving over time. In short, current cyber security policies, rather than being ineffective, are most likely actually helping the situation to a not insignificant degree.

Policy Recommendations

Several policy implications follow from the main finding of the chapter. One cardinal mistake would be to assume that because the security of cyberspace seems to be improving, individuals, companies and governments do not need to act to protect themselves. If perceptions of the security of cyberspace are truly guided by the absolute figures (showing a poor and often worsening environment), then the real improvement in the security of cyberspace is probably driven in part by users' actions intended to counter this dangerous environment by increasing their IT security. More efforts along these lines are needed.

The following recommendations follow from the conclusions of this chapter and the improvement of IT security more generally:

Focus on the individual. The weak point in most IT security systems is often the individual user and not the technical system itself. Spam and phishing emails are designed to capitalize on this weakness, which stems largely from a combination of a lack of knowledge and a likely moral hazard to do with individual responsibility for cyber security. In other words, many people do not know enough to not click the link in a phishing email and many people will likely click links in a work environment that they would never click at home because there is an IT staff to deal with the consequences and individual accountability for data breaches is inconsistent.

Detect and counter new vulnerabilities faster by relying on open source software where possible. Open source software, such as SSL, is often more secure than strictly proprietary programs because it can be examined by so many eyes, although some examples, such as the

Heartbleed exploit in SSL, show that all software is vulnerable. Many, indeed most, individuals with a computer science or computer engineering background are committed to ideals of an open and free Internet. For that to occur, programs need to be secure. Available open source software tends to get examined more often because it is publicly available and this reveals vulnerabilities faster, leading to quicker fixes and more security. In comparison, proprietary programs tend, in general terms, to eventually get leaked, so criminals have access to that code too, but it is not examined by as many eyes, leading to less security on average (Clarke, Dorwin and Nash, n.d.).

Reduce the ability of state security agencies to retain zero-day exploits for law enforcement or national security purposes by requiring that they be disclosed to the software developer within a reasonable timeframe.¹⁹

The US National Security Agency's (NSA's) policy toward zero-day vulnerabilities is one example of the problem of retention by state agencies. According to government sources, the NSA apparently must tell a company that it has discovered a zero-day exploit in its system (Zetter 2014). The major caveat to this requirement is that the NSA can closely guard its knowledge of the zero-day exploit if national security or law enforcement needs dictate (ibid.). Many, if not most, computer programs can be used the world over, so a zero-day exploit in nearly any program can theoretically have national security or law enforcement purposes because it could be used by adversaries of the United States. In the interregnum, while governments sit on zero-day exploits waiting for the chance to use them, the vulnerabilities can also be discovered by criminal elements and used to launch cyber attacks. Creating stricter rules around the disclosure of zero-day exploits, likely along the lines of a reasonable time frame for retention, perhaps on the order of six months to one year after discovery, would help limit the use of these exploits for criminal purposes.

Develop international agreements on spam, phishing emails and other forms of web-based attacks. Some agreements, particularly to do with spam, already exist. As the Internet spreads globally, the reach of these agreements must also spread. Bringing new nations into the potential agreements is also needed. In the case of some attacks, such as DDoS attacks, no agreement exists and there is much more to be done. Figuring out uniform rules to govern these different forms of cyber attack is an important step going forward.

Figure out ways — either through market mechanisms, state intervention or some combination of both — to spread out the costs of cybercrime. As shown above, the Internet contributes a lot more to the global economy than is taken away due to the costs of cybercrime. Overall average organizational cost due to data breaches, for example, is

¹⁹ I am grateful to Melissa E. Hathaway for suggesting this framing of this recommendation.

only a few thousand dollars for every billion dollars of global GDP that the Internet generates. At a global level, the costs of cybercrime are negligible, when you see how much the Internet is contributing to global GDP. Yet, these costs can cause individual firms considerable hardship. Cybercrime insurance is the likely way forward. In this vein, market mechanisms can help protect firms from the costs of cybercrime via a market-driven pricing mechanism that focuses on the risk and potential damage of a cyber attack. Governments could also intervene in the market to regulate the cost of cybercrime insurance and potentially even provide insurance themselves to help protect firms, possibly using a social, rather than a market, discount rate. In all likelihood, a combination of both market and state involvement in the insurance market is needed, especially in the short run, as the market is new and rife with imperfect information. The core idea is that some of the tremendous wealth generated by the Internet should be allocated toward insuring that the actual firms affected by data breaches are not completely destroyed by cybercrime.

Private companies whose operations rely on the Internet need to do more to protect themselves through training, capacity building and investment in IT security systems, at times supported by government grants in the case of SMEs.

The choice of who to target for a cybercrime is likely to be driven by two factors: the probability of successfully targeting the company and the size of the prize to be had.²⁰ Large companies tend to invest more in absolute terms in IT security than SMEs, making them more secure. At the same time, larger companies also offer a more tantalizing target than SMEs as they have more to steal. SMEs, in contrast, tend to invest less in IT security, making them easier targets, but are a less alluring prize for cybercriminals due to their smaller size. Essentially, all businesses are vulnerable. An important secondary implication is that rigorous efforts to provide for IT security at one level can actually displace criminals to another part of the economy, so if larger companies respond to insecurity in cyberspace with large investments in IT security, SMEs might be targeted more frequently. Recognizing this, there is a place for a government grant system to help SMEs develop better IT security so that they are not targeted disproportionately by cybercriminals.

²⁰ Another way to express this notion is that the probability of success ($p = 0$ to 1) discounts the value of what can be taken via a cyber attack ($X = 0$ through ∞). The basic cybercrime equation becomes $P(X)$. For example, a cyber attack that is 50 percent likely to succeed and that is targeting a prize worth, say, \$1,000,000 results in \$500,000 worth of prospective benefit ($0.50[1,000,000] = 500,000$). Likewise, a cybercrime that was 100 percent likely to succeed, but which the prize was only worth \$500,000, would also be worth a total of \$500,000 to the cybercriminal. In short, the difficulty of the attack and the size of the prize both matter when a cybercriminal is picking a company to target.

Norton Symantec, Kaspersky Lab and other cyber security companies should start to collect and represent their data on cybercrime in normalized terms rather than as absolute or year-over-year figures. Understanding the level of insecurity that exists in cyberspace is vitally important and should form the basis of all public and corporate policy going forward. To get an accurate picture of the situation, the numbers on new vectors of attack, web-based attacks and the costs of cybercrime all need to be normalized around the growing size of cyberspace, otherwise a false impression is given, as shown in this chapter. Norton Symantec, Kaspersky Lab and other companies of this sort could help provide valuable data for policy makers by developing — and publicly sharing — clear normalized numbers.

This chapter has shown that the security of cyberspace is actually greater than the impression one gets when looking at the commonly used absolute figures. When the vectors of cyber attack, the occurrence of cyber attacks and the cost of data breaches are normalized around the growing size of cyberspace, the situation seems much less grim.

Acknowledgements

This chapter has benefited from a number of capable eyes. Vivian Moser and Carol Bonnett of CIGI's publications team strengthened the language immensely. And, in no particular order, Andy Wyckoff, Simon Palamar, Laura DeNardis, Melissa Hathaway, Fen Osler Hampson, Gordon Smith, Bill Graham, David Clark and Leanna Ireland all provided terrific comments on the substance and style of the chapter. Their efforts made it far stronger and sharpened the analysis and ideas. The remaining errors are mine and mine alone.

WORKS CITED

- BBC News. 2015. "Cyber War Games to be Staged By UK and US." BBC News, January 16. www.bbc.com/news/uk-politics-30842669.
- CIGI-Ipsos. 2014. "Global Survey on Internet Security and Trust." www.cigionline.org/internet-survey.
- Cisco Systems. 2009. "Cisco Visual Networking Index: Forecast and Methodology, 2008–2013." www.cisco.com/web/BR/assets/docs/whitepaper_VNI_06_09.pdf.
- . 2010. "Cisco Visual Networking Index: Forecast and Methodology, 2009–2014." http://large.stanford.edu/courses/2010/ph240/abdul-kafi1/docs/whitepaper_c11-481360.pdf.
- Clarke, Russel, David Dorwin and Rob Nash. n.d. "Is Open Source Software More Secure?" http://courses.cs.washington.edu/courses/csep590/05au/whitepaper_turnin/oss%2810%29.pdf.
- Dean, David, Sebastian Digrande, Dominic Field, Andreas Lundmark, James O'Day, John Pineda and Paul Zwillenberg. 2012. *The Connected World: The Internet Economy in the G20: The 4.2 Trillion Growth Opportunity*. Boston Consulting Group. www.bcg.com/documents/file100409.pdf.
- Europol. 2015. "Botnet Taken Down through International Law Enforcement Cooperation." Europol, February 25.
- Evans, Dave. 2011. "The Internet of Things: How the Next Evolution of the Internet is Changing Everything." Cisco White Paper. April. www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf.
- Finkle, Jim, Soham Chatterjee and Lehar Maan. 2014. "EBay Asks 145 Million Users to Change Passwords after Cyber Attack." Reuters, May 21. www.reuters.com/article/2014/05/21/us-ebay-password-idUSBREA4K0B420140521.
- Gandalf Group. 2014. "The 36th Quarterly C-Suite Survey: Cyber-Security, Trade Agreements and Foreign Investment." www.gandalfgroup.ca/downloads/2014/C-Suite%20Presentation%20Q3%202014%20Oct%2020%20TC.pdf.
- Gow, Brad. n.d. "Data Security Breaches: More Reach and Frequency Requires More Diligence." Zurich. www.zurich.com/NR/rdonlyres/C4FC10D0-2156-42F8-84E7-63C3BF69B6B6/0/Tech_Cold2_DataBreach.pdf.

- Griffith-Jones, Stephany, Eric Helleiner and Ngaire Woods. 2010. *The Financial Stability Board: An Effective Fourth Pillar of Global Economic Governance*. Special Report. Waterloo, ON: CIGI.
- International Telecommunication Union. 2015. Statistics. www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx.
- Internet Live Stats. 2015a. "Internet Users." www.internetlivestats.com/internet-users/.
- . 2015b. "Total Number of Websites." www.internetlivestats.com/total-number-of-websites/#trend.
- Kaspersky Lab. 2008. "Kaspersky Security Bulletin 2008." <http://securelist.com/analysis/kaspersky-security-bulletin/36241/kaspersky-security-bulletin-statistics-2008/>.
- . 2009. "Kaspersky Security Bulletin 2009." <http://securelist.com/analysis/kaspersky-security-bulletin/36284/kaspersky-security-bulletin-2009-statistics-2009/>.
- . 2010. "Kaspersky Security Bulletin 2010." <http://securelist.com/analysis/kaspersky-security-bulletin/36345/kaspersky-security-bulletin-2010-statistics-2010/>.
- . 2011. "Kaspersky Security Bulletin 2011." <http://securelist.com/analysis/kaspersky-security-bulletin/36344/kaspersky-security-bulletin-statistics-2011/>.
- . 2012. "Kaspersky Security Bulletin 2012." <http://securelist.com/analysis/kaspersky-security-bulletin/36703/kaspersky-security-bulletin-2012-the-overall-statistics-for-2012/>.
- . 2013. "Kaspersky Security Bulletin 2013." http://media.kaspersky.com/pdf/KSB_2013_EN.pdf.
- . 2014. "Kaspersky Security Bulletin 2014." <http://cdn.securelist.com/files/2014/12/Kaspersky-Security-Bulletin-2014-EN.pdf>.
- Norton Symantec. 2009. *Internet Security Threat Reports*. http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiv_04-2009.en-us.pdf.
- . 2010. *Internet Security Threat Reports*. http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xv_04-2010.en-us.pdf.
- . 2011. *Internet Security Threat Reports*. www4.symantec.com/mktginfo/downloads/21182883_GA_REPORT_ISTR_Main-Report_04-11_HI-RES.pdf.
- . 2012. *Internet Security Threat Reports*. www.trustico.com/news/internet_security_reports/internet_security_report_2012_04.en-us.pdf.
- . 2013. *Internet Security Threat Reports*. www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v18_2012_21291018.en-us.pdf.
- . 2014. *Internet Security Threat Reports*. www.symantec.com/content/en/us/enterprise/other_resources/b-istr_main_report_v19_21291018.en-us.pdf.
- . 2015. *Internet Security Threat Reports*. www4.symantec.com/mktginfo/whitepaper/ISTR/21347932_GA-internet-security-threat-report-volume-20-2015-social_v2.pdf.
- O'Toole, James. 2014. "Mobile Apps Overtake Desktop Internet Usage in U.S." CNN Money, February 28. <http://money.cnn.com/2014/02/28/technology/mobile/mobile-apps-internet/>.
- Pélessié du Rausas, Matthieu, James Manyika, Eric Hazan, Jacques Bughin, Michael Chui and Rémi Said. 2011. *Internet Matters: The Net's Sweeping Impact on Growth, Jobs and Prosperity*. McKinsey & Company. May. www.mckinsey.com/insights/high_tech_telecoms_internet/internet_matters.
- Ponemon Institute. 2011. *2011 Cost of Data Breach Study: United States*. Ponemon Institute Research Report. www.ponemon.org/local/upload/file/2011_US_CODB_FINAL_5.pdf.
- . 2013. *2013 Cost of Data Breach Study: A Global Analysis*. Ponemon Institute Research Report. www4.symantec.com/mktginfo/whitepaper/053013_GL_NA_WP_Ponemon-2013-Cost-of-a-Data-Breach-Report_daiNA_cta72382.pdf.
- . 2014. *2014 Cost of Data Breach Study: Global Analysis*. Ponemon Institute Research Report. www-935.ibm.com/services/us/en/it-services/security-services/cost-of-data-breach/.
- Radicati Group. 2013. "Email Market 2013–2017." www.radicati.com/wp/wp-content/uploads/2013/11/Email-Market-2013-2017-Executive-Summary.pdf.

- Richwine, Lisa. 2014. "Sony's Hacking Scandal Could Cost the Company \$100 Million." *Business Insider*, December 9. www.businessinsider.com/sonys-hacking-scandal-could-cost-the-company-100-million-2014-12.
- Royal Pingdom. 2009. "Internet 2008 in Numbers." <http://royal.pingdom.com/2009/01/22/internet-2008-in-numbers/>.
- . 2010. "Internet 2009 in Numbers." <http://royal.pingdom.com/2010/01/22/internet-2009-in-numbers/>.
- . 2011. "Internet 2010 in Numbers." <http://royal.pingdom.com/2011/01/12/internet-2010-in-numbers/>.
- . 2012. "Internet 2011 in Numbers." <http://royal.pingdom.com/2012/01/17/internet-2011-in-numbers/>.
- . 2013. "Internet 2012 in Numbers." <http://royal.pingdom.com/2013/01/16/internet-2012-in-numbers/>.
- Silver-Greenberg, Jessica, Matthew Goldstein and Nicole Perloth. 2014. "JPMorgan Chase Hacking Affects 76 Million Households." *The New York Times*, October 2. http://dealbook.nytimes.com/2014/10/02/jpmorgan-discovers-further-cyber-security-issues/?_r=0.
- Statista. 2015. "Number of Smartphones Sold to End Users Worldwide from 2007 to 2014 (in million units)." www.statista.com/statistics/263437/global-smartphone-sales-to-end-users-since-2007/.
- Statistics Brain. 2015. "Google Annual Search Statistics." www.statisticbrain.com/google-searches/.
- StatsCounter. 2015. "Top Desktop, Console and Tablet Browsers per Country, Oct 2014." <http://gs.statcounter.com/#browser-ww-monthly-201410-201410-map>.
- Verisign. 2009. *Domain Name Industry Brief 6* (1). Verisign. www.verisigninc.com/assets/domain-name-report-feb09.pdf.
- . 2010. *Domain Name Industry Brief 7* (1). Verisign. www.verisigninc.com/assets/domain-name-report-feb10.pdf.
- . 2011. *Domain Name Industry Brief 8* (1). Verisign. www.verisigninc.com/assets/domain-name-report-feb-2011.pdf.
- . 2012. *Domain Name Industry Brief 9* (1). Verisign. www.verisigninc.com/assets/domain-name-brief-march2012.pdf.
- . 2013. *Domain Name Industry Brief 10* (1). Verisign. www.verisigninc.com/assets/domain-name-brief-april2013.pdf.
- . 2014. *Domain Name Industry Brief 11* (1). Verisign. www.verisigninc.com/assets/domain-name-report-april2014.pdf.
- . 2015. *Domain Name Industry Brief 12* (1). Verisign. www.verisigninc.com/assets/domain-name-report-march2015.pdf.
- Woodcock, Bill and Vijay Adhikari. 2011. "Survey of Characteristics of Internet Carrier Interconnection Agreements." Packet Clearing House. May 2. www.pch.net/resources/papers/peering-survey/PCH-Peering-Survey-2011.pdf.
- Zetter, Kim. 2014. "U.S. Gov Insists it Doesn't Stockpile Zero-Day Exploits to Hack Enemies." *Wired*, November 17. www.wired.com/2014/11/michael-daniel-no-zero-day-stockpile/.

ABOUT THE AUTHOR

Eric Jardine joined CIGI as a research fellow in May 2014 in the Global Security & Politics Program. He contributes to CIGI's work on Internet governance, including the CIGI-Chatham House-sponsored Global Commission on Internet Governance. His current research focuses on cyber security, cyberterrorism, cybercrime and cyber protest. He holds a Ph.D. in international relations from the Norman Paterson School of International Affairs at Carleton University.

CHAPTER TWO: THE IMPACT OF THE DARK WEB ON INTERNET GOVERNANCE AND CYBER SECURITY

Michael Chertoff and Toby Simon

Copyright © 2014 by the Centre for International Governance Innovation
and the Royal Institute of International Affairs

INTRODUCTION

In his advance in the Battle of the Persian Gate in 331 BC, Alexander the Great passed into the Persian Gate with little or no resistance. Convinced that he would not encounter enemy forces, Alexander neglected to send scouts ahead, and thus walked into a Persian ambush while crossing a pass on his way to Persepolis. Persian troops on either side rained boulders and arrows down on the invaders. The Macedonians suffered heavy casualties, losing entire platoons, and were forced to withdraw. Alexander then gathered intelligence from a local shepherd to encircle the Persian army in a pincer attack. His knowledge of the larger terrain helped him to outflank the Persians and emerge victorious.

Four hundred years later, seven Roman legions, some 44,000 men, marched into the searing Mesopotamian desert. They had come to the eastern province of the kingdom of Parthia seeking conquest and plunder, but, caught unaware by the uncharted terrain, the legions were almost annihilated. Most of the Romans were either slaughtered or captured and enslaved. Their commander was decapitated, and his head was used as an ornament at the banquet of the Parthian king. The Battle of Carrhae was a disaster almost unmatched in the otherwise glorious history of the Roman army. Twenty thousand were killed and 10,000 taken prisoner. It was the worst Roman defeat since the dreadful loss to Hannibal at Cannae in 216 BC. It was the result of engaging an unknown adversary, in an unknown land.

These two anecdotes remind us of the importance of reconnaissance, and the need to better understand what is beneath the surface. The deep and the dark Web can pose unseen threats. About 40 percent of the world's population uses the Web for news, entertainment, communication and myriad other purposes (International Telecommunication Union 2014). As more and more people become Internet users, they are actually finding less of the data that is stored online. Only a sliver of what we know as the World Wide Web is easily accessible.

The surface Web, which people use routinely, consists of data that search engines can find and then offer up in response to queries. This is only the tip of the iceberg — a traditional search engine sees about 0.03 percent of the information that is available (Bergman 2001). Much of the rest is submerged in what is called the deep Web. Also known as the “Undernet,” “invisible Web” and the “hidden Web,” it consists of data that cannot be located with a simple Google search.

In order to formulate comprehensive strategies and policies for governing the Internet, it is important to consider insights on its farthest reaches — the deep Web and, more importantly, the dark Web. This chapter endeavours to provide a broader understanding of the dark Web and its impact on our lives.

CONTEXT

On November 3, 2014, the newly appointed director of Britain's Government Communications Headquarters, Robert Hannigan, warned that US tech giants such as Twitter, Facebook and WhatsApp have become the “command-and-control networks of choice for terrorist and criminals” (Hannigan 2014). Hannigan's statements were among the most critical of American technology firms by the head of a major intelligence agency and, more significantly, a close ally. The accusation went beyond what US officials have said so far about Apple, Google and others that are now moving toward sophisticated encryption of more and more data on phones and email systems (Wilber 2014).

This revelation was closely followed by a low-profile post by Facebook informing users that it is now hosted directly on the Tor network (Lee 2014). The Tor link — <https://facebookcorewwi.onion/> — was described more as an experiment by the company, to enable it to learn over time by providing an onion address¹ for Facebook's mobile website. Incidentally, Facebook is the first US tech giant to provide official support for Tor, a network built to allow citizens to surf the Web without being tracked and publish content that would not show up in normal search engines.

Hannigan's understanding of how the coupling of social media and the dark Web could create extremely powerful, encrypted, decentralized and anonymous propaganda networks for terrorist organizations may be what prompted him to speak out. The recent surge in the number of European nationals sympathetic to or actively supporting organizations like ISIL (the Islamic State of Iraq and the Levant) or al-Qaeda in Syria and Iraq is definitely a huge cause of worry for Western democracies. Social media platforms have proven themselves valuable recruitment tools for campaigns of all types. It is of little surprise, then, that in recent years, terrorist groups such as al-Qaeda and ISIL have successfully employed Twitter to recruit volunteers and be active in supporting their cause (Coughlin 2014). The intent is clearly to “humanize” the movement and reach broader audiences.

Beyond propaganda, cyberspace allows groups to spread particular knowledge in new and innovative ways. The kinds of tools that allow social organizations such as the Khan Academy to help kids around the world learn math and science have also given terrorist groups unprecedented ways to discuss and disseminate tactics, techniques and procedures. Recipes for explosives are readily available on the Internet and terror groups have used the Internet to share

¹ An onion address designates an anonymous hidden service reachable via the Tor network.

designs for improvised explosive devices instantly across conflict zones from Syria to Afghanistan (Singer 2012).

The visible side of the Internet includes sites that can be found through an ordinary search, while the invisible side—the deep Web—includes sites or networks that cannot be accessed by regular means. This includes databases, academic journals, private networks and so on. Most of the content located in the deep Web exists in websites that require a search that is not implicitly illicit. However, an intensive search will find the dark Web. The dark Web is a small portion of the deep Web that has been intentionally hidden.

Deep Web Resources:

- Dynamic content
- Unlinked content
- Private Web
- Contextual Web
- Limited access content

Accessing the Deep Web:

- Custom Web crawlers use key terms provided by users or collected from the query interfaces to query a Web form and crawl the deep Web resources.
- Commercial search engines have begun exploring alternative methods to crawl the deep Web. The Sitemap Protocol (first developed and introduced by Google in 2005) is a mechanism that allows search engines and other interested parties to discover deep Web resources on particular Web servers (Google 2014).

While innovative methods have been developed for monitoring content on the visible Web in recent years, there are almost no similar tools for the dark Web. Providing evidence showing that the dark Web has turned into a major platform for global terrorism and criminal activities is crucial in order for the necessary tools to be developed for monitoring all parts of the Internet.

THE INTERNET, THE WORLD WIDE WEB AND THE DEEP WEB

Many people use the terms Internet and World Wide Web interchangeably, but in fact the two terms are not synonymous. The Internet and the Web are two separate but related things.

The Internet is a massive network of networks — a networking infrastructure. It connects millions of computers together globally, forming a network in which any computer can communicate with any other computer, as long as they are both connected to the Internet.

On the other hand, the World Wide Web, or simply the Web, is a way of accessing information over the medium of the Internet. It is an information-sharing model that is built on top of the Internet. The Web uses the Hypertext Transfer Protocol, only one of the languages spoken over the Internet, to transmit data. The Internet, not the Web, is also used for email, which relies on Simple Mail Transfer Protocol, Usenet news groups, instant messaging and File Transfer Protocol. The Web, therefore, is just a portion of the Internet, albeit a large one (Beal 2010). Finally, the deep Web is, put simply, the part of the Web that is hidden from view. It is World Wide Web content that is not part of the surface Web. It cannot be accessed by normal search engines. This massive subsection of the Internet is more than 500 times bigger than the visible Web (Barker and Barker 2013).

THE DARK WEB

The dark Web is the portion of the deep Web that has been intentionally hidden and is inaccessible through standard Web browsers. Dark Web sites serve as a platform for Internet users for whom anonymity is essential, since they not only provide protection from unauthorized users, but also usually include encryption to prevent monitoring.

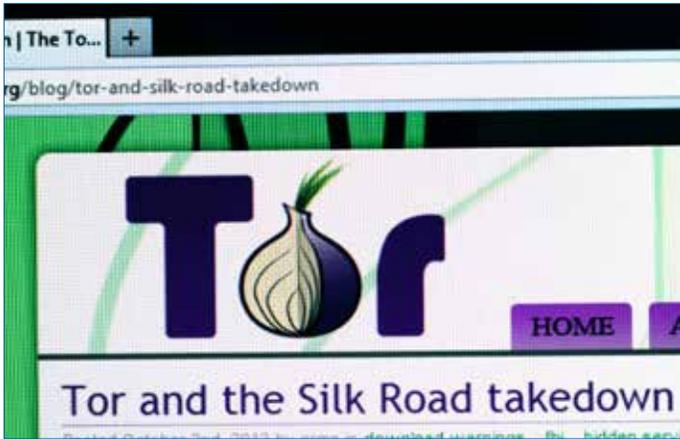
A relatively known source for content that resides on the dark Web is found in the Tor network. The Tor network is an anonymous network that can only be accessed with a special Web browser, called the Tor browser (Tor 2014a). First debuted as The Onion Routing (Tor) project in 2002 by the US Naval Research Laboratory, it was a method for communicating online anonymously. Another network, I2P, provides many of the same features that Tor does. However, I2P was designed to be a network within the Internet, with traffic staying contained in its borders. Tor provides better anonymous access to the open Internet and I2P provides a more robust and reliable “network within the network” (Tchabe and Xu 2014).

Usage

The ability to traverse the Internet with complete anonymity nurtures a platform ripe for what are considered illegal activities in some countries, including:

- controlled substance marketplaces;
- credit card fraud and identity theft; and
- leaks of sensitive information.

Silk Road was an online marketplace that dealt with contraband drugs, narcotics and weapons. In 2013, the US Federal Bureau of Investigation (FBI) shut down the website. But like the mythical Hydra, the website resurrected as Silk Road 2.0 within a month. It took the FBI another year to track down its administrator and servers (Mac 2014).



A Tor Project website blog page discussing the takedown of Silk Road (an online marketplace that dealt with contraband drugs, narcotics and weapons) by the FBI. iStock.

It should also be noted that Tor empowers anyone who wants control over his or her online footprint. The positive value of such a tool is huge for some groups, such as whistle-blowers who report news that companies would prefer to suppress, human rights workers struggling against repressive governments and parents trying to create a safe way for their children to explore the Web.

Defining Attributes

Anonymity, from the Greek word *anonymia*, refers to the state where one's personal identity is not publicly known. Each day, our Web actions leave footprints by depositing personal data on the Internet. This information composes our digital identity — our representation in cyberspace.

Internet anonymity is guaranteed when Internet Protocol (IP) addresses cannot be tracked. Tor client software routes Internet traffic through a worldwide volunteer network of servers, hiding user's information and eluding any activities of monitoring. This makes the dark Web very appropriate for cybercriminals, who are constantly trying to hide their tracks (Paganini 2012).

The dark Web is also the preferred channel for governments to exchange documents secretly, for journalists to bypass censorship of several states and for dissidents to avoid the control of authoritarian regimes (Gehl 2014). Anonymous communications have an important place in our political and social discourse. Many individuals wish to hide their identities due to concerns about political or economic retribution.

Onion routing is a technique for anonymous communication over a computer network. Messages are repeatedly encrypted and then sent through several network nodes, called onion routers. Like someone peeling an onion, each onion router removes a layer of encryption to uncover routing instructions, and sends the message

to the next router, where the process is repeated. This technique prevents intermediary nodes from knowing the origin, destination and contents of the message (Tor 2014a).

CYBERCRIME IN THE DARK WEB

Peter Grabosky (2001) notes that virtual crime is not any different than crime in the real world — it is just executed in a new medium: “‘Virtual criminality’ is basically the same as the terrestrial crime with which we are familiar. To be sure, some of the manifestations are new. But a great deal of crime committed with or against computers differs only in terms of the medium. While the technology of implementation, and particularly its efficiency, may be without precedent, the crime is fundamentally familiar. It is less a question of something completely different than a recognizable crime committed in a completely different way.”

Drugs, Weapons and Exotic Animals

Websites such as Silk Road act as anonymous marketplaces selling everything from tame items such as books and clothes, to more illicit goods such as drugs and weapons. Aesthetically, these sites appear like any number of shopping websites, with a short description of the goods, and an accompanying photograph (Bartlett 2014).

Stolen Goods and Information

It is correct to assume that dedicated sites facilitate users to trade in both physical and proprietary information, including passwords and access to passwords for surface Web paid-pornography sites and PayPal passwords (Westin 2014). PayPal Store, Creditcards for All and (Yet) Another Porn Exchange are active websites that offer such services.

Murder

The Assassination Market website is a prediction market where a party can place a bet on the date of death of a given individual, and collect a payoff if the date is “guessed” accurately. This incentivizes the assassination of individuals because the assassin, knowing when the action will take place, could profit by making an accurate bet on the time of the subject's death. Because the payoff is for knowing the date rather than performing the action of the assassination, it is substantially more difficult to assign criminal liability for the assassination (Greenberg 2013). There are also websites to hire an assassin — popular ones are White Wolves and C'thuthlu (Pocock 2014).

Terrorism

The dark Web and terrorists seem to complement each other — the latter need an anonymous network that is

readily available yet generally inaccessible. It would be hard for terrorists to keep up a presence on the surface Web because of the ease with which their sites could be shut down and, more importantly, tracked back to the original poster.

While the dark Web may lack the broad appeal that is available on the surface Web, the hidden ecosystem is conducive for propaganda, recruitment, financing and planning, which relates to our original understanding of the dark Web as an unregulated space.

Hacktivism

More radical critics and hacktivists occupy part of the political dissidence space. The group Anonymous, commonly associated with Occupy Wall Street and other cyber activism, is one prominent hacktivist group (Jones 2011).

Exploit Markets

Exploits are malware based on software's vulnerabilities — before they are patched. Zero-day exploits target zero-day vulnerabilities — those for which no official patch has been released by the vendor. "Zero-day" refers to the fact that the programmer has had zero days to fix the flaw. Exploit markets serve as platforms for buying and selling zero-day exploits, and an exploit's price factors in how widely the target software is used as well as the difficulty of cracking it (Miller 2007).

One of the things driving the rapid rise in cybercrime is that the cybercriminal does not have to be a master hacker since the exploits can be bought.

— Sir David Bruce Omand, GCB

Illegal Financial Transactions

Websites such as Banker & Co. and InstaCard facilitate untraceable financial transactions through various methods. They either launder bitcoins by disguising the true origin of the transactions or give users an anonymous debit card issued by a bank. Users are also given virtual credit cards issued by trusted operators in the dark Web (Dean 2014).

Buying stolen credit card information has never been easier. A website called Atlantic Carding offers this service, and the more you pay, the more you get. Up for grabs are business credit card accounts and even infinite credit card accounts associated with ultra-high-net-worth individuals. The user's details — name, address and so on — are available at an additional cost (Dahl 2014).

The Hidden Wiki

The main directory on the dark Web is the Hidden Wiki. It also promotes money laundering services, contract killing, cyber attacks and restricted chemicals, along with instructions to make explosives. As with other dark Web sites, the links to these sites frequently change to evade detection (Williams 2011).

Human Experimentation

The Human Experiment was a website that detailed medical experiments claimed to have been performed on homeless people who were usually unregistered citizens. According to the website, they were picked up off the street, experimented on and then usually died. The website has been inactive since 2011 (Falconer 2012).

Heist

There are many rob-to-order pages available in the dark Web, hosted by people who are good at stealing and will steal anything that you cannot afford or just do not want to pay for (Siddiqui 2014).

Arms Trafficking

Euroarms is a website that sells all kinds of weapons that can be delivered to your doorstep anywhere in Europe. The ammunition for these weapons is sold separately — that website has to be tracked down separately on the dark Web (Love 2013).

Gambling

Many popular bitcoin gambling sites block US IPs because they are afraid of prosecution from the United States, which has a tight hand on gambling in the United States. With the help of the dark Web, users of these sites can continue gambling by disguising their US IP (O'Neill 2013).

Pedophilia

Pedophilia, or CP (for child pornography) as it is commonly referred to on the dark Web, is extremely accessible. Pornography is accepted on the surface Web with some regulation. The dark Web offers various types of sites and forums for those wishing to engage in pedophilia (Greenberg 2014).

THE CASE FOR ONLINE ANONYMITY

Like any technology, anonymity can be used for both good and bad purposes. Many people do not want the things they say online to be connected with their offline identities. They may be concerned about political or economic retribution, harassment or even threats to their lives. Instead of using their true names to communicate, these people choose to speak using pseudonyms or anonymously. Listed below are

a few scenarios where users turn to the online anonymity provided by Tor (Tor 2014b).

Civilians

- protection of privacy from unscrupulous marketers and identity thieves
- protection of communications from irresponsible corporations
- protection of children online
- to research sensitive topics
- to circumvent censorship

Militaries

- field agents
- hidden services of command and control
- intelligence gathering

Journalists and Their Audience

- to help Internet users in countries without safe access to free media
- to write about local events to encourage social change and political reform
- to avoid risking the personal consequences of intellectual curiosity

Law Enforcement

- online surveillance
- sting operations
- maintaining anonymous tip lines

Activists and Whistle-blowers

- to report abuses from danger zones
- anonymous blogging
- to speak out about government corruption

For these individuals and the organizations that support them, secure anonymity is critical. It may literally save lives. While the undesired effects of Tor must be recognized, the complexities and varied situations should make us suspicious of sweeping imperatives. Policies should be crafted to specific contexts (Marx 1999).

MONITORING THE DARK WEB

The dark Web, in general, and the Tor network, in particular, offer a secure platform for cybercriminals to support a vast amount of illegal activities — from anonymous marketplaces to secure means of communication, to an untraceable and difficult to shut down infrastructure for deploying malware and botnets.

As such, it has become increasingly important for security agencies to track and monitor the activities in the dark Web, focusing today on Tor networks, but possibly extending to other technologies in the near future.

Due to its intricate webbing and design, monitoring the dark Web will continue to pose significant challenges. Efforts to address it should be focused on the areas discussed below (Ciancaglini et al. 2013).

Mapping the Hidden Services Directory

Both Tor and I2P use a domain database built on a distributed system known as a “distributed hash table,” or DHT. A DHT works by having nodes in the system collaboratively take responsibility for storing and maintaining a subset of the database, which is in the form of a key-value store. Due to the distributed nature of the hidden services domain resolution, it is possible to deploy nodes in the DHT to monitor requests coming from a given domain.

Customer Data Monitoring

Security agencies could benefit from analyzing customer Web data to look for connections to non-standard domains. Depending on the level of Web usage at the customer side, this may not help in tracking down links to the dark Web, but it may still provide insights on activities hosted with rogue top-level domains. This can be done without intruding on the user’s privacy as only the destinations of the Web requests need to be monitored and not who is connecting to them.

Social Site Monitoring

Sites such as Pastebin are often used to exchange contact information and addresses for new hidden services. These sites would need to be kept under constant observation to spot message exchanges containing new dark Web domains.

Hidden Service Monitoring

Most hidden services to date tend to be highly volatile and go offline very often, coming back online later under a new domain name. It is essential to get a snapshot of every new site as soon as it is spotted, for later analysis or to monitor its online activity. While crawling the clear Internet is

usually an operation involving the retrieval of resources related to a site, this is not recommended in the dark Web. There is the possibility of automatically downloading content such as child pornography, the simple possession of which is considered illegal in most countries.

Semantic Analysis

Once the data for a hidden service (any of the websites on the dark Web) is retrieved, building a semantic database that contains important information about a hidden site can help track future illegal activities on the site and associate them with malicious actors.

Marketplace Profiling

Finally, it would be helpful to focus on profiling transactions made on dark Web marketplaces to gather information about sellers, users and the kinds of goods exchanged. Individual profiles could be built up over time.

CONCLUSION

The deep Web — in particular, networks on the dark Web such as Tor — represents a viable way for malicious actors to exchange goods, legally or illegally, in an anonymous fashion.

The lack of observable activities in unconventional dark Web networks does not necessarily mean they do not exist. In fact, in agreement with the principle that inspires the dark Web, the activities are simply more difficult to spot and observe. A driving factor for the marketplace is critical mass. Operators in the dark Web are unlikely to need a high level of stealth unless the consequences, if they are discovered, are sufficiently severe. It is conceivable that sites may come online at specific times, have a brief window of trading, then disappear, making them more difficult to investigate.

Recent revelations about wide-scale nation-state monitoring of the Internet and recent arrests of cybercriminals behind sites hosted in the dark Web are starting to lead to other changes. It would not be surprising to see the criminal underbelly becoming more fragmented into alternative dark nets or private networks, further complicating the job of investigators.

The dark Web has the potential to host an increasingly large number of malicious services and activities and, unfortunately, it will not be long before new large marketplaces emerge. Security researchers have to remain vigilant and find new ways to spot upcoming malicious services to deal with new phenomena as quickly as possible.

WORKS CITED

- Barker, Donald I. and Melissa Barker. 2013. *Internet Research Illustrated*. Independence, KY: Cengage Learning, C-4.
- Bartlett, Jamie. 2014. "Dark Net Markets: The eBay of Drug Dealing," *The Observer*, October 5.
- Beal, Vangie. 2010. "The Difference between the Internet and World Wide Web." Webopedia, June 24.
- Bergman, Michael K. 2001. "White Paper: The Deep Web: Surfacing Hidden Value." <http://quod.lib.umich.edu/j/jep/3336451.0007.104?view=text;rgn=main>.
- Ciancaglini, Vincenzo, Marco Balduzzi, Max Goncharov and Robert McArdle. 2013. "Deepweb and Cybercrime: It's Not All About TOR." Trend Micro Research Paper. October.
- Coughlin, Con. 2014. "How Social Media Is Helping Islamic State to Spread Its Poison." *The Telegraph*, November 5.
- Dahl, Julia. 2014. "Identity Theft Ensnarers Millions while the Law Plays Catch Up." CBS News, July 14.
- Dean, Matt. 2014. "Digital Currencies Fueling Crime on the Dark Side of the Internet." Fox Business, December 18.
- Falconer, Joel. 2012. "A Journey into the Dark Corners of the Deep Web." *The Next Web*, October 8.
- Gehl, Robert W. 2014. "Power/Freedom on the Dark Web: A Digital Ethnography of the Dark Web Social Network." *New Media & Society*, October 15. <http://nms.sagepub.com/content/early/2014/10/16/1461444814554900.full#ref-38>.
- Google. 2014. "Learn about Sitemaps." [ps://support.google.com/webmasters/answer/156184?hl=en](https://support.google.com/webmasters/answer/156184?hl=en).
- Grabosky, Peter. 2001. "Virtual Criminality: Old Wine in New Bottles?" *Social & Legal Studies* 10: 243-49. <http://sls.sagepub.com/content/10/2/243.full.pdf>.
- Greenberg, Andy. 2013. "Meet the 'Assassination Market' Creator Who's Crowdfunding Murder with Bitcoins." *Forbes*, November 18.
- . 2014. "Over 80 Percent of Dark-Web Visits Relate to Pedophilia, Study Finds." *Wired*, December 30.
- Hannigan, Robert. 2014. "The Web Is a Terrorist's Command-and-Control Network of Choice." *The Financial Times*, November 3.
- International Telecommunication Union. 2014. "The World in 2014: ICT Facts and Figures." www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2014-e.pdf.

- Jones, Melanie. 2011. "Anonymous NYSE: Wall Street 'Hacktivism' Exposes Downside of Anonymity." *International Business Times*, October 11.
- Lee, David. 2014. "Facebook Sets Up 'Dark Web' Link to Access Network via Tor." BBC News, November 3.
- Love, Dylan. 2013. "There's a Secret Internet for Drug Dealers, Assassins, and Pedophiles." *Business Insider*, March 6.
- Mac, Ryan. 2014. "Feds Shutter Illegal Drug Marketplace Silk Road 2.0, Arrest 26-Year-Old San Francisco Programmer." *Forbes*, November 6.
- Marx, Gary T. 1999. "What's in a Name? Some Reflections on the Sociology of Anonymity." <http://web.mit.edu/gtmarx/www/anon.html>.
- Miller, Charlie. 2007. "The Legitimate Vulnerability Market: Inside the Secretive World of 0-day Exploit Sales." <http://weis2007.econinfosec.org/papers/29.pdf>.
- O'Neill, Patrick Howell. 2013. "Inside the Bustling, Dicey World of Bitcoin Gambling." *The Daily Dot*, December 17.
- Paganini, Pierluigi. 2012. "The Good and the Bad of the Deep Web." *Security Affairs*, September 17.
- Pocock, Zane. 2014. "How to Navigate the Deep Web." *Critic*, Issue 03, March 19.
- Siddiqui, Sameer Iqbal. 2014. "Real Power of Deep Web and How to Harness It." *Real Hackers Point* (blog), June 19.
- Singer, Peter W. 2012. "The Cyber Terror Bogeyman." November. The Brookings Institution. www.brookings.edu/research/articles/2012/11/cyber-terror-singer.
- Tchabe, Gildas Nya and Yinhua Xu. 2014. "Anonymous Communications: A Survey on I2P." www.cdc.informatik.tudarmstadt.de/fileadmin/user_upload/Group_CDC/Documents/Lehre/SS13/Seminar/CPS/cps2014_submission_4.pdf.
- Tor Project. 2014a. "Tor: Overview." www.torproject.org/about/overview.html.en.
- . 2014b. "Inception." www.torproject.org/about/torusers.html.en.
- Westin, Ken. 2014. "Stolen Credit Cards and the Black Market: How the Deep Web Underground Economy Works." *LinkedIn*, August 22.
- Wilber, Del Quentin. 2014. "U.S. Law Enforcement Seeks to Halt Apple-Google Encryption of Mobile Data." *Bloomberg News*, September 30.
- Williams, Christopher. 2011. "The Hidden Wiki: An Internet Underworld of Child Abuse." *The Daily Telegraph*, October 27.

ABOUT THE AUTHORS

Michael Chertoff, chairman and co-founder of the Chertoff Group and senior of counsel, Covington & Burling LLP, was secretary of the US Department of Homeland Security from 2005 to 2009. Previously, he was a US Court of Appeals judge and chief of the US Department of Justice Criminal Division. He is a magna cum laude graduate of both Harvard Law School and College. He is a commissioner with the Global Commission on Internet Governance.

The Chertoff Group is a global security advisory firm that provides consulting, business development and through Chertoff Capital, merger and acquisition advisory services for clients in the security, defence and government services industries. The Chertoff Group also advises public and private enterprises on their own physical and cyber security. With decades of trusted leadership experience across both government and financial services, the Chertoff Group advises clients on how to manage their risk, protect against a broad array of threats and crises, and grow their business within a complex security market.

Tobby Simon is president of The Synergia Foundation, an applied research think tank that works closely with academia, industry and polity to establish impactful solutions in the areas of geo-economics and geo-security. Tobby is a commissioner at the Global Commission for Internet Governance and an advisory board member of the Centre for New American Security. He is a graduate of the Harvard Business School and a research affiliate at the Massachusetts Institute of Technology. He is currently pursuing his Ph.D. at the National Institute of Advanced Studies in Bangalore, India.

The Synergia Foundation strives to help, contribute and influence public policy, private initiatives and international relations in making our region a better place for its present and future inhabitants. The foundation has over 400 years of combined experience in strategic thinking, and has multi-disciplinary teams that pursue high quality non-partisan research and draw on its global network of resources to offer the most comprehensive research analysis and impactful solutions.

**CHAPTER THREE:
THE DARK WEB DILEMMA:
TOR, ANONYMITY AND ONLINE POLICING**

Eric Jardine

Copyright © 2014 by the Centre for International Governance Innovation
and the Royal Institute of International Affairs

INTRODUCTION

The dark Net¹: its very name brings to mind images of shadowy alleys, malicious, hard-faced individuals and socially damaging activity. The dark Net is a part of the Internet that most people probably do not know how to access, nor want to explore. Only a special web browser is needed to reach it.² One such browser, embedded in a larger networked system, is The Onion Router (TOR) network.³

A lot happens via Tor. This chapter runs through some of what goes on in the dark Net, with a particular focus upon how the anonymity of the widely used Tor browser allows for both nefarious and noble undertakings. It uses evidence from a variety of news accounts and secondary literature to detail how anonymity can be used as a tool of those that want to undertake socially damaging activity. It also uses the results of a recently conducted study on Tor usage rates that shows empirically that people in politically repressive countries are often driven to use the anonymity network out of necessity (Jardine, n.d.).

The basic story to emerge from all of this evidence is that an anonymity-granting system such as Tor, as with other technologies, is just a tool. Like fire, a hammer or a car, the Tor network can both improve life and provide the means to take it away. What matters is not what the technology is, but how it is used and what the net effect turns out to be.

Framed from this perspective, the focus of public debate should move away from demonizing the technology, or looking for quick technological fixes, toward the idea that, like every other aspect of human society, the dark Net needs to be policed. This recommendation is particularly relevant for liberal democratic countries, where the dark side of anonymity imposes the highest costs and the benefits of Tor are least pronounced. Ideally, policing needs to be undertaken within clearly defined, rule-based limits. That is no different than the rest of society. Sometimes, as the saying goes, the more things change, the more they stay the same.

1 The dark Web and the dark Net are used interchangeably throughout this chapter and mean the same thing.

2 The borders of the dark Web are blurry. See, for example, Chertoff and Simon (2015). For the purposes of this chapter, the dark Web can be defined as a part of the Internet that is only possible because of online anonymity. This definition does not imply that online anonymity is enough to create the dark Web, only that the dark Web cannot exist without it. In social science terms, online anonymity is a necessary, but not a sufficient, cause of the dark Web.

3 The Tor browser is the entry point of focus to the dark Web for this chapter, but there are other ways into the Internet's underground. The Tor browser is also one of the main gateways to anonymity in this chapter. Again, others exist.

The next section describes the Tor-hosted dark Net. Following that, the chapter discusses the negative effects generated by the anonymity of the dark Web. The third section presents new statistical evidence to show that sometimes the anonymous network is used for good. The fourth discusses the policy implications that flow from the dual nature of the technology, in particular, how online policing of the dark Web has proven to be just as effective as offline policing. The only way forward is to police the dark Web, just as we police all aspects of society.

TOR AND THE DARK NET

Under normal circumstances, when you are trying to access the Web, you send a signal from your device across the Internet to the server that hosts the material that you want to view. That can be a cat meme, a pornographic video, a news organization's webpage or whatever else might tickle your fancy. The server then returns the data to your device. The relationship is direct. Your request is sent via the networks of the Internet to the place that holds the information you want to view and it is sent back.

Because of this directness, our Internet service providers (ISPs) know our names, addresses, search histories and the sites that we are visiting. It is also how the websites we view know our unique Internet Protocol (IP) address. It is because of this direct connection that companies such as Amazon know everything we view and even how long we have lingered upon a page. Law enforcement agencies are able to capitalize on this directness and can pinpoint who posted what information on an online chat forum.

Tor accesses information on the Web in much the same way, but it breaks up the direct connection. After a fashion, the Tor browser is a bit like an anonymous version of the children's game of telephone. You send your request for a particular video or bit of information to a computer somewhere in the Tor network. This computer then relays that information on to another computer somewhere else in the network. Once again, this computer simply relays your request onward to yet another machine. This third machine in the game of telephone then requests the information you want to view and sends it back to you along a similar, disjointed path.

Breaking up the request in this way means that different people can see different parts of what you are viewing online, but it is exceptionally difficult, although not impossible, for any one person to connect all the dots to pinpoint who you actually are (Owen and Savage 2015). Your ISP, for example, which normally knows exactly what sites you are visiting, can only see that you are sending a request to the first computer in the network. On the other end of things, the website can tell a lot about the computer that is accessing their content, but this information does not relate to your computer, instead linking to the last of the three computers in the game of telephone. The computers

in the relay system know about their neighbour, but no more than that. The first link knows you and the middle computer, but not the end computer or the content viewed. The middle link knows the first computer and the end computer, but not you or the destination of your request. The end computer knows the destination and the middle computer, but not who you are. Layered onto this broken routing of your request is the heavily encrypted signal that prevents data flowing across the Tor network from being accessible to prying eyes.

Tor is not just a way to view online content anonymously. You can also host content, but only in a way that is accessible to other users of the Tor browser. Put another way, you can be the one running the website to which people venture for their bits of information, whatever they might be. The process by which anonymity is obtained is similar to that laid out above. The website itself moves around from server to server in the Tor network. Changes to the website are made using the same three-relay system that is used to prevent the website or server from knowing who is hosting the page. Anonymity is secured.

Finally, it is important to clear up at the outset that, while large parts of the dark Web are only reachable via Tor, the Tor browser itself can actually be used for other far more innocent purposes, such as simply surfing the day-to-day Web, free from the constraints of censored content and concern over state or corporate surveillance. If you try to download and use Tor (a process that is very easy), you will find that you never need to venture into the seedy underbelly of the Internet if you do not want to. Instead, you can use the Tor browser just like Google Chrome or Mozilla Firefox to check news websites, look at funny memes or anything else you would do normally when browsing the Internet.⁴ Even these routine activities are rendered anonymous by Tor.

The end result of this system is a way to use the Internet anonymously, with all the immunity that provides. Clearly, as shown in the next section, that anonymity opens the door to abuses.

THE DARK SIDE OF ONLINE ANONYMITY

The dark Net certainly is the seedy underbelly of the Internet. Its sordid nature is exemplified in a few stories about drugs, assassination, trolling and child abuse.

In the early years of this decade, a site popped up on the dark Web called Silk Road. The reference to the ancient trading route from the Orient to Europe was not a mistake. The website was like an illegal version of Amazon, eBay,

⁴ Plug-ins are limited on Tor, so you might not have the full range of functionality you would on another web browser, but the idea that you could just use Tor for your normal Internet activity is valid.

Kijiji or Craigslist. It aimed to connect sellers of items ranging from drugs to assassinations-for-hire with eager customers with money to burn.

Silk Road started in February 2011. One study observed activity on the website during a six-month period in 2012 and found that Silk Road, while selling all sorts of illegal content, was mostly a proverbial “drugstore.” Categorizing all the things that were for sale on the site, the authors found that “the four most popular categories are all linked to drugs,” along with 90 percent of the top 10 categories and 80 percent of the top 20 (Christin 2012, 8). The transactions were anonymous due to the use of the Tor network and payments were made with a so-called cryptocurrency known as bitcoin, which is a purely digital means of payment that leaves no trace.

Silk Road quickly surpassed other illegal market sites, with its revenue and traffic expanding rapidly. In an uncomfortable mix of metaphors, the site was owned by a then 29-year-old man who went by the moniker Dread Pirate Roberts — taken straight out of the 1980s movie *The Princess Bride*. By 2012, the site operators were earning upwards of \$92,000⁵ per month, as people were flocking to the site to buy and sell items on the illegal market. The audacity of Silk Road’s illegal activities led US Senator Charles Schumer to call for the site to be shut down in June 2011, noting that it is “more brazen than anything else by light-years” (cited in Koebler 2012).

The investigation into Silk Road started in 2011, when an informant broke word of activity on the illegal marketplace site to personnel at the Department of Homeland Security. Operation “Marco Polo,” as the investigation came to be called, quickly expanded to encompass personnel from the Federal Bureau of Investigation (FBI), DHS, Drug Enforcement Administration, Internal Revenue Service and others (Zetter 2013).

As the law enforcement net was closing in on the Dread Pirate Roberts, the modern-day bandit got desperate, even offering \$80,000 to an undercover agent to assassinate a former site administrator who had been captured by the police and turned state’s evidence. The police staged the killing of the site administrator just to draw the noose that much tighter around Dread Pirate Roberts’ neck (ibid.). Ross Ulbricht, the Dread Pirate Roberts, was arrested in October 2013 and the site was taken down. It was a clear victory.

It was also very short-lived. Silk Road 2.0 popped up on the dark Net in November 2013, just one month after the arrest of Ulbricht. Again, the website expanded rapidly, quickly having as many as 150,000 active users and processing, according to FBI records, as much as \$8 million in monthly sales (Cook 2014). Within a year, this new incarnation of the illegal marketplace was taken down

⁵ All currency in this chapter is in US dollars.

and Blake Benthall, the Silk Road 2.0 site administrator and former Space-X employee, was arrested.

Another win, another drop in the pond. Silk Road 3.0 was online within a few hours of Benthall's arrest (Knibbs 2014). The cycle goes on, like a globe-spanning game of whack-a-mole.

The dark recesses of the dark Web are also populated with proverbial trolls, some of whom use Tor to maintain their anonymity, some of whom do not. We have all come across Internet trolls. They surf the Web, posting inflammatory comments, aiming for nothing more than to wreck someone's day, often just for the fun of it.

Consider this telling story of trolling and a needlessly ruined life on the 4chan /b/ board (Bartlett 2014, 13–19).⁶ A young university student named Sarah ventured half-naked via a posted photograph into the chat board filled with dark Web trolls. Her first photo spawned a number of requests for further nudity, which she willingly provided. The requests built gradually to a terrible point. One request asked her to pose naked with her name written on her body. She did it. Another request asked her to pose naked with any medications that she might be taking. She did that, too.

From there, the situation got really ugly. Her mistake was providing the trolls of the dark Web with enough information to identify her. They found her school, accessed its directory and got her full name, address, phone number and other contact information. Facebook searches revealed her social media profile. From there, the anonymous chatters of the /b/ chatroom then began a “doxing”⁷ campaign to wreck her reputation by sharing her naked photos with everyone she had even a slight connection with. Why? Because they could. The viciousness of it all needs to be recounted verbatim to be believed:

Anonymous: “she gave her first name, her physician's full name, and even the dormitory area she lives in[.] [S]he wants to be found” (Bartlett 2014, 15).

⁶ 4chan actually forbids users from posting using Tor or a virtual private network (VPN) to hide their true identities, so this example might seem slightly outside of the scope of the chapter. It is, nevertheless, included for a couple of reasons. First, the extent to which the ban on Tor is followed or enforceable is quite unclear, and it is likely that many routinely violate it. Additionally, the nature of the 4chan board itself provides a degree of anonymity to posters, with users actually being told not to use any identifiable information in their profiles. So, even if the operators would (assuming the rule prohibiting Tor is followed) be able to backtrace posts to a particular person if law enforcement requested it, the ability of people to behave badly because of the anonymity of the board is still present.

⁷ Doxing basically involves taking people's personal information and spreading it as widely as possible.

Anonymous: “here is a list of all her Facebook friends. You can message friends, and all their own friends, so that anyone with a slight connection to sarah [sic] via friend of friend knows” (ibid., 17–18).

Anonymous: “so has somebody started messaging her friends or family or can I begin with it? (ibid., 18).

Anonymous: “[xxxxx] is her Fone [sic] number — confirmed” (ibid.).

Anonymous: “just called her, she is crying. She sounded like a sad[,] sad sobbing whale” (ibid.).

Anonymous: “Is anyone else continually calling?” (ibid.).

The attacks were personal, devastating and brutal. But the anonymous posters of the 4chan /b/ board were also remorseless.

Anonymous: “If [she] was clever she would have g[ot] t[he] f[***] o[ut][,] she didn[']t, therefore she deserves the consequences” (ibid., 19).

Anonymous: “I don't give a s*** what happens either. Bitch was camwhoring while she had a boyfriend” (ibid., 19).

The torment promised to be long-lived as well. Amid the maelstrom, Sarah had tried to minimize the damage by deleting her social media accounts, such as Facebook, to limit the trolls' access to the people she knew. But, as one troll noted, the Internet's memory is eternal:

Anonymous: “Eventually once all this settles she will reactivate it [her Facebook account] and she will have her jimmies rustled once more. She will now never know peace from this rustling. And she's going to have one embarrassing f***ing time with her family” (ibid., 16).

It is sad to see even one life wrecked by a couple of bad choices that are then magnified by the destructive behaviour of anonymous trolls. But this case is in no way an isolated incident. One study found that upwards of two-thirds of people between the ages of 13 and 22 have been bullied online (Butterly 2013). And while certainly not all bullying goes on in the dark Web — Facebook being a key vehicle of bullying — some of the most egregious often does. It is widespread, malicious and at times enabled by anonymity-granting tools like Tor. Its consequences are both individually and socially destructive.

With its illegal drug and weapon markets and online trolls, the dark Web seems immoral and unscrupulous, but the scary part is that the shadows of the dark Web can actually get even darker. Nothing makes that point more clearly than the prevalence of child abuse imagery on the dark Net.

In 2011, Europol, coordinating with 13 national governments, launched Operation Rescue. The concerted law enforcement action uncovered 670 suspects and led to 184 arrests on child abuse imagery-related charges (Europol 2011). In July 2014, the UK's National Crime Agency arrested some 650 people on various child abuse charges, ranging from the possession of images to the actual abuse of minors (BBC 2014a). In 2015, another 50 suspects were identified in Northern Ireland and 37 charges were laid (BBC 2015). These are just a few examples of the successful instances of law enforcement uncovering pedophilia rings in the recesses of the dark Web.

Unfortunately, as Gareth Owen and Nick Savage (2015) point out in their study for the Global Commission on Internet Governance, the problem of child abuse images on the dark Web is probably even more widespread than the record of arrests would lead us to believe. In their innovative study, Owen and Savage actually volunteered a couple of servers to host the Tor network at their university in Portsmouth, United Kingdom. Over a period of several months, they categorized the type of websites found on the Tor-hosted dark Web. They found that the available sites ranged from whistle-blower chatrooms to pornography sites, illegal markets and child abuse sites. This last category accounted for only a small fraction of all sites hosted on the dark Web. Unfortunately, they also found that over 80 percent of the actual traffic along the Tor anonymity network went to this small proportion of sites (*ibid.*).

The lesson from all this is that anonymity allows the dark Web to be a very nasty place indeed, and Tor makes this type of behaviour possible. Illegal markets selling drugs and guns to whomever will pay, malicious trolls and those who want to harm children, are but a few of the villainous activities going on within the lower recesses of the Internet.

The Virtuous Protection of the Shadows

But the anonymity of the technology of Tor cuts both ways — while people can use the network for villainous purposes, people can also use it for good.

Anonymity is important for the possibility of democracy. Anonymity provides space for people to think and voice opinions that are against the grain. Anonymity ensures both protection for an individual that holds a minority point of view and a window of opportunity for the majority consensus to be challenged by outside ways of thinking. As noted in a US Supreme Court decision, *McIntyre v. Ohio*

Elections Commission, “Anonymity is a shield from the tyranny of the majority.... It thus exemplifies the purpose behind the Bill of Rights and of the First Amendment in particular: to protect unpopular individuals from retaliation...at the hand of an intolerant society” (cited in Electronic Frontiers Foundation, n.d.). Without a healthy public debate encompassing all viewpoints, democracy shrivels. In non-democratic countries, the presence of anonymity is the only way that people can voice contrary points of view against despotic regimes in the hope of securing political freedom.

For its part, the Tor Project website maintains that political activists, reformers, journalists, civil rights workers and development workers can use Tor in repressive countries to circumvent censorship and, to some extent, avoid the prying eyes of state and corporate surveillance. Use of the anonymity network has also been suggested by human rights groups. Reporters Without Borders, for example, recommends the use of Tor as a part of its journalist's “survival kit” (Murray 2014). In its somewhat older report on Internet usage in China, *Race to the Bottom*, Human Rights Watch supported the use of Tor (Human Rights Watch 2006). And the human rights advocacy group, Global Voices, suggests that Tor is useful for dissidents and activists (Global Voices, n.d.).

All of these suggestions for using the Tor network might or might not translate into people actually using it for noble purposes in regimes that mean harm to ordinary citizens. Unlike the high-profile instances of online drug busts and child pornography arrests, which are both on the moral high ground and newsworthy, there are few public stories of political activists using Tor. Repressive regimes do not broadcast when they break the encryption of the network and throw people who are simply asserting their right to free expression into dank prisons. Those who use Tor to avoid surveillance or to circumvent censorship are also not likely to publicly proclaim the specifics of their use of the network (or even that they use it at all), since the whole point of the system is to keep one's online activity anonymous.

There is another way, however, to discover whether people really do use the Tor dark Web in repressive countries. Rather than have people self-report that they use the network, one can look at usage numbers per country. While many of the specifics are unknowable (as befits an anonymity network), the Tor Project provides data on the number of users of its network per country. Of course, each country has a different number of Internet users and different rates of Internet penetration, so it is not just a matter of counting the number of users and saying that the largest number of users are in either repressive or liberal regimes. Instead, to get at whether the level of political rights in a country drives usage of the Tor network, you need to use special statistical methods with a large sample size that can account for other factors that might also lead

to people using the network. The process is not as complex as it sounds. At their most basic level, statistical methods can give you an impression of how often a certain level of political rights is associated with either high or low use of the Tor bridge network, given the effect of a host of other factors.

Before turning to the outcome of the statistical tests, it is a good idea to explain how statistical methods can produce some relatively intelligible answers. The basic question explored in another study (Jardine, n.d.) is whether people used Tor more as political repression increased from 2011 to 2013, which gets at the problem of whether the anonymity of the dark Net can actually provide a cloak to protect those that want to exercise their rights to free speech and freedom of information.

On the one side, the political rights measure used in this study ranges along a scale from one to seven, and is taken from a widely used measure known as the Freedom in the World Index (Freedom House 2015). The index is scored like a game of golf: lower is better. A score of one, in this case, is the best, and a seven is the worst. Liberal democratic countries such as Canada, the United States and the United Kingdom score a one on the political rights index. Highly repressive countries such as Chad and Swaziland score a seven. The rest of the countries of the world are spread between these extremes.

The outcome to be explained is the use of the Tor network in different countries per year, with a specific focus on the use of what are known as bridges. Tor bridges are another name for the relay computers in the game of anonymous telephone. The one distinction is that unlike normal relays, bridges are not listed publicly, which makes them a better tool for people to circumvent censorship and surveillance in repressive regimes.

Since you would expect more people to use Tor (or for that matter anything) in a large population compared to a small one, the numbers for the outcome to be explained are expressed as a rate per 100,000 Internet users per year in a country. A simple example can demonstrate why normalizing the data in this manner is important: in 2013, the United States had 147,207 Tor bridge users, while Canada only had 23,795 users. On the face of it, it seems like Americans use the network a lot more than Canadians — and in one sense they do. But, as a population as a whole, America actually uses the network less. The United States has 55 Tor bridge users per 100,000 Internet users, while Canada has 79 users per 100,000 Internet users. Expressed in these terms, Canadians actually use Tor bridges at a 43.6 percent greater rate than their American cousins. The normalization matters.

Other factors in addition to political rights also drive use of the Tor network. So, to get a realistic picture of the effect of differing level of political rights, those conditions need to

be factored into the equation. Wealth is important to take into consideration because it affects access to information technologies and national bandwidth capabilities. Internet penetration rates are important because someone needs to be able to access the Internet if they are going to actually use Tor. Exposure to foreign ideas and influences also matter, as people need to know about Tor in order to use it in the first place. Education matters because people need to have a certain level of comfort with information and communications technology in order to use something outside the norm such as Tor. Intellectual property rights regimes matter because they can increase the incentive to use Tor to download illegal movies and songs. The statistical tests include all these factors.⁸

Putting all these numbers to use and running some statistical regressions shows a clear relationship between Tor bridge use per 100,000 Internet users per year and a country's level of political rights. And while political rights do matter, they also *do not* matter in a straightforward way. Rather than use of the Tor network simply increasing as the political rights situation worsens in a country, the relationship between rights and the use of Tor is shaped like a "U." In other words, political rights tend to drive usage rates the most in both highly liberal countries such as Canada and highly illiberal countries such as Swaziland.

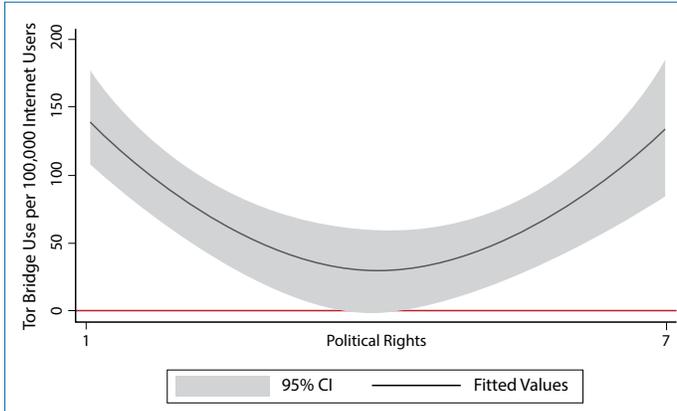
Figure 1 shows how the relationship unfolds across the actual data. As the political rights situation moves from a country such as the United States (political rights = 1) to a country such as Honduras (political rights = 4), political rights tend to drive use of the Tor network less and less. Beyond that low point, a worsening political rights situation starts to drive people toward using the Tor network again, as evidenced by the right hand side of the U-shaped relationship.

The magnitude of the effect is knowable, too. Table 1 shows on average just how much a change in the level of political rights in a country matters. Moving from a 1 to a 4 on the political rights scale results in a total reduction of 174.99 users of Tor bridges per 100,000 Internet users per year. Going the other way, from a 5 to a 7 on the political rights scale, leads to a total increase of 68.42 Tor bridge users per 100,000 Internet users per year. In short, political rights matter a fair bit for use of the network.

The obvious question at this stage is why do political rights matter in this way? Why form a U-shaped relationship? The reason is that a political regime drives the domestic population's opportunity to use Tor, as well as their need to do so, with the former factor declining as repression increases and the latter rising as political rights decline.

⁸ Issues of multicollinearity are discussed in detail in Jardine (n.d.).

Figure 1: Political Rights and Tor Bridge Usage



Source: Jardine (n.d.).

Table 1: Changing Political Rights and Tor Bridge Use per 100,000 Internet Users per Year

Change in political rights	Change in Tor bridge users per 100,000 Internet users
1 to 2	84.77 less
2 to 3	58.33 less
3 to 4	31.89 less
4 to 5	5.45 less
5 to 6	20.99 more
6 to 7	47.43 more

Source: Jardine (n.d.).

Opportunity, for instance, starts out high in liberal countries, as there are few restrictions on the use of encrypted or anonymous technologies such as Tor. Indeed, a large portion of the Tor Project’s funding comes from the US government and the genesis of the program is in US military research labs. As the level of political rights declines, the opportunity to use the anonymity-granting technology worsens, as repressive regimes throw up roadblocks — for example, legislation and technical blocking mechanisms — to prevent people from using the system. China, for instance, has been fairly successful at blocking Tor (MIT Technology Review 2012). Russia, a six on the political rights scale, has offered \$110,000 to the person or organization that can crack the encryption and anonymity of the Tor network (BBC 2014b).

Opportunity counts for a lot, so, if it is nearly costless to do so, people will use programs such as Tor for illegal reasons, to circumvent censorship and surveillance by both states and corporations, or simply to support the idea that the anonymous use of the Internet should be something that is valued in society. The result of high opportunity is the high use of anonymity-granting technologies in highly liberal countries.

Need, for its part, is low in liberal countries. People do not have to use Tor in order to do their legal online activity in liberal countries with a strong tradition of rights protection, although the extent to which people should take more steps to be anonymous online, even in liberal regimes, is an open question. As the level of political repression goes up within a country, the need to use anonymity-granting programs like Tor rises.

This growing need drives people to use Tor in repressive regimes. Here again, the motives vary. Some will do so for illegal purposes. But others will use the network to blow the whistle on corruption, to freely express their political viewpoints, to circumvent censorship and to avoid direct surveillance of their online activity.⁹

The basic point is that repression and the violation of political rights does drive people to use the anonymity network. Oftentimes, people in repressive regimes simply cannot freely express their points of view, circumvent the censorship of important information or avoid the prying eyes of the state without encrypted and anonymous programs such as Tor. Some of what people do online with Tor in repressive regimes will be innocuous and some will even be illicit or illegal, but much of it will be virtuous and aimed at nothing more than exercising some fundamental political rights.

THE POLICY DILEMMA: A DUAL-USE TECHNOLOGY

As demonstrated, Tor is basically a dual-use technology: it can be used for truly awful purposes as well as for good. How it is used matters most, similar to other tools that humanity has invented. We discovered how to harness fire to keep us warm, but then learned that it can be used to ravage and burn. We discovered steel and now use it to make buildings that touch the sky, but before that we learned it can be used to make swords or guns to take lives. The human story is riddled with the invention of technologies that can be used for both good and ill.

Discussions of the use of the Tor network, like discussions of encryption in general, are highly polarized. The one side asserts that the technology needs to be as close to unbreakable as possible so that nefarious actors cannot gain access. A back door into an encrypted system cannot be given only to law enforcement and somehow kept from criminals and political despots. Once an entryway exists, the system is vulnerable. Indeed, purposeful back doors can lead to less privacy, more vulnerabilities as new systems interact with past software and even make

⁹ Because Tor has distinct encryption, repressive regimes can often tell when someone is using the program, even if they cannot tell what is being done with the system. Paradoxically, this effort to dodge surveillance of content might put an individual under more scrutiny as the use of encrypted technologies raises red flags in many repressive regimes.

governments and service providers tantalizing targets of cybercrime, as they possess the proverbial keys to the kingdom (Abelson et al. 2015).

The other side of the debate asserts that encrypted and anonymous technologies such as Tor hinder law enforcement. FBI Director James B. Comey exemplifies this position. In October 2014, he pointed straight to the other half of the polarization in a speech at the Brookings Institution in Washington, DC:

Encryption isn't just a technical feature; it's a marketing pitch. But it will have very serious consequences for law enforcement and national security agencies at all levels. Sophisticated criminals will come to count on these means of evading detection. It's the equivalent of a closet that can't be opened. A safe that can't be cracked. And my question is, at what cost? (Comey 2014)

Indeed, at what cost? In one way, the policy issue as it specifically relates to the Tor network boils down to a question about whether the technology does more harm than good. What matters is a net assessment of the impact of the technology. There is no straightforward answer to this question, but the evidence presented here suggests a painful underlying truth — how you frame the parameters of the cost-benefit calculus affects the answer you get.

The uncomfortable reality is that liberal democratic nations that developed and host much of the Tor network are actually having to deal with most of the negative consequences of the system while reaping few of the benefits. The opportunity to use the technology in liberal countries means that Silk Road, trolls and anonymous child abuse websites proliferate, but the gains (dodging the prying eyes of state or corporate content surveillance and circumventing censorship) are fairly minimal. Other, less cumbersome programs (private search engines, such as Duck Duck Go, and VPNs) exist and have roughly the same effect as Tor with more download speed and less potential for abuse, as they retain user data and can cooperate with law enforcement if approached with a valid warrant. Therefore, unless people are engaged in outright illegal activities, the need to use a full-blown anonymity program such as Tor in liberal democratic countries is also limited, because of the presence of constitutional and legal protections of citizen rights, although it is important to not under-represent the extent to which the rapidly evolving nature of the technology of the Internet has outpaced the ability of the legal system to deal with new challenges to citizen's fundamental rights. Based upon the evidence presented above, the idea that Tor provides net benefits to society in liberal democratic countries is unlikely. It most likely does more harm than good.

If the frame of reference is shifted to the net costs or benefits of Tor in a highly repressive country, however, the cost-benefit outcome changes radically. Dissidents, journalists, human rights activists and even ordinary citizens in repressive countries all benefit from the Tor network, even if some of these people might use it for nefarious purposes. In the end, the Tor anonymity network in regimes with low political rights is definitely more beneficial overall.

The implicit policy question to come out of this is whether people in liberal countries are willing to pay the cost of the existence of a system such as Tor, given that the benefits are not evenly distributed globally. People in Western countries might decide that the costs are simply not worth it and opt for a state-driven clamp down on the system. This decision would have serious implications for the effectiveness of the Tor network as it functions well in repressive regimes only because most of its infrastructure (computers and servers) reside in liberal countries. Without innumerable volunteered computers around the world, the anonymity of the network would be limited and the ability of Tor to cloak those in need in repressive regimes would be stymied.

WHAT IS TO BE DONE? POLICING

Even if people in liberal countries decide that a program such as Tor is not worth having, the odds of destroying anonymity-granting technologies in general in an era of a global Internet are pretty slim. Tor might be knocked offline, but other programs would simply emerge and take their place. Unless you break the global Internet (which would be excessively expensive in terms of lost GDP), it is simply not possible to prevent people from building technologies that ensure the anonymous use of the Web. In other words, the problem of a dual-use technology like Tor is not likely to go away any time soon. We are stuck with both the good and the bad.

Rather than looking for quick and final fixes (such as destroying Tor outright or altering the technology through back doors in encryption for law enforcement), a more realistic way forward is to focus on actively policing the network.

In the offline world, peace and order are maintained in every segment of society through judicious policing. Socially destructive behaviours are deemed illegal. Crimes are recorded. And criminals are arrested, prosecuted and sent to jail. It is actually ridiculous to think that as more of our daily lives and activities shift online, the online world would not also need to see a rapid expansion of policing efforts to accommodate the shift in our attention and activity.

There has already been some movement in this direction by police forces around the world (Omand, forthcoming). This movement shows that online policing of the dark Web

is in fact possible, expedient and often at least as effective as offline policing.

Despite the use of the Tor network to host the various Silk Road illegal marketplaces, for example, the owners and operators of the sites — as well as many of the largest sellers — were identified and arrested. These arrests show the effectiveness of online policing. The takedown of Silk Road 1.0 is instructive. Police caught the Dread Pirate Roberts through a combination of technological means and the double-edged sword of online anonymity.

Tor is obviously a technically heavy system. And technology played a role in the capture of the server hosting the Silk Road and the ultimate arrest of Ross Ulbricht. In the initial prosecution filing against Ulbricht, the FBI indicated that it found the location of the Silk Road server in Iceland due to a misconfiguration on the illegal market's log-in page, which allowed investigators to type in "miscellaneous" characters in a CAPTCHA window that returned IP address information.¹⁰ Upon further snooping, the FBI realized that the IP address provided by the log-in page did not correspond to a known node in the Tor network, and was likely the actual physical address of Silk Road rather than a relay in the system (Greenburg 2014a). Technology is a fickle mistress and it betrayed those that were relying upon it to do harm.

Of course, others doubt whether the characters typed into the CAPTCHA by the FBI were really miscellaneous, charging instead that they were actually lines of code designed to hack the log-in page by duping it into thinking the entries were actually administrative commands (Greenburg 2014b). Both accounts are plausible. Silk Road 2.0, for example, was not vulnerable to the same flaw, suggesting either that Silk Road 1.0 was taken down by a configuration issue or perhaps by a now-patched vulnerability (Brandom 2015). Indeed, Ulbricht's defence during his trial that there was an illegal search due to how the FBI found the Silk Road server fell apart. He was sentenced to more than life in prison (Thielman 2015).

Silk Road 1.0 was also taken down because of the very thing that allowed it to operate in the first place: anonymity. Anonymity, that core feature provided by the Tor browser, does not stop law enforcement. Instead, it actually makes law enforcement efforts, in some ways, easier. Buyers or sellers on Silk Road, trolls and child abusers cannot say for sure who they are dealing with in an online world. Anonymity limits attribution, but it cuts both ways. No further evidence is needed than the Dread Pirate Roberts, who offered money to an undercover cop to undertake an assassination of a former site administrator. Child abuse sites are also routinely infiltrated by law enforcement.

Police from the United Kingdom and Australia, for example, infiltrated one online child abuse ring of up to 70,000 members "to identify the members who posed the greatest danger to children. Police also sometimes posed as children online as part of the investigation" (NBC News, n.d.).

Online policing is also as expedient as offline policing. The anonymity of Tor does not necessarily slow down law enforcement efforts. The fact that the Silk Road networks were taken down, often within a year of their launches, shows the speed at which online policing can work. As a parallel analog example, Project DISTRESS was launched in Manitoba, Canada, in October 2013, and culminated 15 months later in the arrest of 14 suspects in a major drug trafficking ring (RCMP 2014). The scope of this real-world effort is smaller than Operation Marco Polo to take down Silk Road 1.0, but the timelines are roughly the same. If anything, the online version was a larger endeavour but took less time to complete. Online policing seems to be at least as quick as its analog cousin.

The fact that new Silk Road marketplaces, trolls or child abuse sites keep popping up in the wake of arrests and shutdowns is also nothing new, and should not be taken as evidence that online policing is not effective. Offline, the arrest of a street-corner drug dealer often leaves a void that is quickly filled by someone else. This does not mean that we should stop arresting drug dealers. It means that we are stuck with the problem of people selling drugs, at least until the demand for what is being sold goes away or the arrest and prosecution for such activity is certain. The same logic applies online. Yes, new sites will always pop up as the old ones are taken down and arrests are made, but this just means that governments need to keep policing the network. It is part of the cost of the Internet. To obtain all the benefits that the Internet provides, we need to ensure it is as safe as possible, but we do not want to destroy it completely, which is the only way prevent crime from occurring online.

The call for greater online policing is not the same as saying the state should be allowed to intervene indiscriminately into people's lives. Offline, the police cannot go into people's homes whenever they want, but they can patrol the streets and catch people in the act of committing crimes. The same sort of logic should apply online. Police should not be allowed to access the data on a person's computer or their ISP records without a warrant. At the same time, they are allowed to sit in chatrooms to monitor conversations and even pose as potential victims to catch predators. They are also allowed to pose as sellers or buyers on illegal markets to track down people who are actually committing crimes. In short, the new "beat" is shifting from the street to the websites and chatrooms of the Internet. This is the reality of the digital age. Certain tactics remain off limits — and law enforcement should not purposefully take advantage of the presence of legal

¹⁰ CAPTCHAs are those website windows with blurry letters and numbers that are designed to fool spamming machines, but allow humans to access a site.

ambiguity to overreach — but the Internet will not work as a global free-for-all.

This policing should also avoid politicizing the core infrastructure of the Internet. As Samantha Bradshaw and Laura DeNardis (n.d.) note, attempting to police intellectual property rights regimes, for example, through the core infrastructure of the Internet (in their case, the Domain Name System) can lead to unintended consequences that risk damaging or even breaking the network. Instead, policing of the dark Web should occur largely on top of the infrastructure at the social or content level. Law enforcement officers should have a presence inside an online chatroom frequented by pedophiles, but they should not manipulate the infrastructure that supports the creation of online chatrooms in the first place.

There is a bit of a tension between the legitimate use of technological methods to identify those that are breaking the law and the idea that manipulating core infrastructure should be off limits. The use of technology to fight crime falls along a continuum. At one end are legitimate technical investigations, such as the methods used to take down Silk Road 1.0. This kind of activity is acceptable because it exploited a weakness in a particular site, rather than trying to break the whole system. At the other end, trying to simply knock Tor offline is a more fundamental politicization of the infrastructure of the system, affecting both the good and the bad indiscriminately, and therefore should be disallowed.

At the margin, there is a lot of ambiguity about what is acceptable. The takedown of Silk Road 2.0 points out the blurry line. To identify the users of Silk Road 2.0, the FBI volunteered “reliable IP addresses” to the Tor hidden services network upon which the newest incarnation of the illegal marketplace was based. This allowed the FBI to subtly change the coding so that they could pinpoint the identity of users that had employed their relays to reach the illegal marketplace. The operators of Tor noted this trick after six months, and provided a patch that once again improved the anonymity of Tor. For Silk Road 2.0 and Blake Benthall, it was too late. The FBI had tracked down the server and 78 sellers and buyers (Brandom 2015). Exploiting the voluntarist nature of the Tor infrastructure is right at the line of unacceptable use of core infrastructure for policing. It was an indiscriminate attack on all Tor users, so it probably went a bridge too far. Either way, the Silk Road 2.0 example highlights the tension.

LIMITATIONS TO ONLINE POLICING AND AREAS FOR POLICY INTERVENTION

There are limits to the effectiveness of online policing that concerted policy actions can help to overcome.

One limitation is that online criminals can be global, even while most law enforcement agencies (Interpol excepted) are local. If a criminal is not in the same jurisdiction as the police that identify his or her actions as illegal, policing gets immensely more complicated. The problem is even more pronounced when Tor bounces your signal around the world, effectively involving multiple jurisdictions. In some cases, policies are in place to allow states to cooperate by sharing evidence across borders. Foremost among these mechanisms are what is known as mutual legal assistance treaties (MLATs).

The problem is the MLAT process is in massive need of reform. Proposals exist for how it should be reformed. One study maintains that MLAT reform must emphasize proportionality, the protection of human rights, transparency, heightened efficiency and scalability if they are to become an effective tool in the international police officer’s tool kit (Woods 2015). That would be a good start.

MLAT reform can certainly help to make the process of Internet policing more effective, but it will not solve the root of the problem, as online crime is highly mobile and can drift to countries that are outside of the effective MLAT regime. For MLATs to work, two states need to have an agreement in place and both need to view something as illegal in order for the process to be effective. Cooperation through the MLAT process is quite likely between liberal democratic countries because they share legal principles and political dispositions. Cooperation on cybercrime is less likely between Western countries and nations such as China and Russia, which disagree on so many fundamental issues. Moreover, at the end of the day, MLAT reform might fail as the Internet governance system is becoming increasingly contentious (Bradshaw et al. 2015). This is not a small hurdle, but it is not insurmountable either.

Other specific efforts at international coordination of law enforcement agencies can do nothing but help. Interpol’s Global Complex for Innovation is a prime example. It aims to build relationships between police forces, increase various countries’ understanding of digital security issues and facilitate capacity building to overcome the fact that many local and national police forces just do not have the resources, training and wherewithal to deal well with cybercrime. More international coordination should help with the trans-border portion of the cybercrime problem.

But coordination failures are not just a problem between nations. Most countries have internal layers of police,

ranging from the national to the local. Coordination failures between these levels can often stymie effective efforts at policing cybercrime. Local and national police have both critical resources and deficiencies in the battle against cybercrime. Local police can often be the first to learn of a cybercrime (say, identity theft or cyber harassment), but often lack the capacity and jurisdiction to act effectively.¹¹ National law enforcement usually has the capacity and jurisdiction to act effectively, but can lack knowledge that a particular cybercrime is occurring.

The strengths and weaknesses of local and national-level law enforcement are complementary. By working together, the knowledge of local police can be paired with the resources and capacity of national law enforcement. Specialization remains efficiency-enhancing here, so local police should not be trying to bust international online fraud rings and national-level law enforcement should not be trying to get local victims to report crimes directly to them (although national-level crime reporting is increasingly effective at scale). Each level should stick to its strengths, but work together in a coordinated way to limit online crime.

Even with greater coordination, more training and capacity are still needed. Local law enforcement, in particular, tends to be undertrained and under-resourced to deal with cybercrime. As Darrel Stephens, executive director of the Major City Chiefs Police Association, noted in 2013, “Most local police do not have the capacity to investigate these cases even if they have jurisdiction” (cited in Sullivan 2013). Stephens is also cognizant of how local police departments will need to adapt, stating further that, “Police will need to become more equipped to deal with cybercrime in the future” (ibid.). And that “most major cities have a limited capability, but more will be required” (ibid.). Many crimes are shifting online, so resources that are otherwise dedicated to policing offline crime could be usefully moved to combat online crime instead. Even with the redistribution of efforts, more resources are needed to effectively combat online crime.

Obtaining more resources at the local level is likely to come with some growing pains. More resources typically follow greater need, but local police face a perverse incentive when it comes to something as foundational to crime fighting as recording that a crime has even occurred. A physical burglary or violent crime in a jurisdiction will faithfully be recorded accurately and quickly in most cases. A cybercrime of harassment or theft is far less likely to be counted. The reason is that it is harder for local police to address these crimes, given resources, capacity and the jurisdiction in which they work. As a result, these crimes

are more likely to remain off the books.¹² To include them would inflate the crime rate in an area and probably the unsolved crime rate as well, all of which reflects poorly upon the local police department.

However, by trying to avoid a rising crime rate, local law enforcement is hamstrung in their ability to solicit or collect new resources or capacity over the long term. Heads might roll if the crime rate goes up in the short run, but this could be a window of opportunity for local police departments that need more training and resources to combat cybercrime. In most cases, a growing need (higher crime rates) is matched with more resources. In the long term, the only way to strengthen local police departments to help them fight cybercrime is to recognize that cybercrime has local victims, even if perpetrators could be anywhere in the world and the jurisdictional lines are blurry.

Increasingly, coordination must also occur between governments and private sector actors. One example of this coordination in action is the recent breakup of a large botnet by European law enforcement and Microsoft (Microsoft News Center 2013). Private companies own and operate much of the software, hardware and networks of the Internet, while law enforcement has the jurisdiction to pursue criminals. Public-private partnerships between law enforcement and private companies will likely be the way of the future. When done well, public-private collaboration can be a massive force multiplier, leading to the more effective policing of the dark Web.

Policing anonymity-granting technologies is also challenging because the system is decentralized, based upon volunteered servers and does not retain data. The messaging application Wickr is an analogous example. They will readily comply with warrants that require access to their servers; however, since they do not retain any data generated by the users of their service, law enforcement cannot find any useful information by searching the system. Tor is similar in that it does not retain data. Additionally, the volunteered nature of the network means that even if someone were logging traffic through Tor relays (which the system is not designed to do), law enforcement in any one country would be hard pressed to find this data. Changing the legal rules so that companies and organizations such as Tor would be required to retain data for a period of time — for instance, six months — would be one way to allow for semi-anonymous communications, but ensure that when law enforcement is cued to a potential crime, they can get access to what they need. The big problem with this approach is how it would be applied in repressive regimes. In those countries, even a six-month retention of data can lead to imprisonment for activists, journalists and human

11 Many countries have national-level information collection agencies, so information about ongoing crimes is not always clustered at the local level. This varies by country and likely by crime type as well.

12 At the 2015 Global Conference on Cyberspace in The Hague, Richard Clayton pointed out that this happens. To the extent that I may have misunderstood his point, the fault is my own.

rights workers. As a result, those behind Tor would never accept a mandated retention period of data.

A final limitation is that cybercrime is rapidly increasing, which threatens to overwhelm any and all available policing capacity of nations. Cybercrime is certainly going up, but it is not as bad as we commonly think it is. The key reason is that cyberspace is actually growing as fast, and sometimes faster, as the growth in new vulnerabilities, web-based attacks and the costs of cybercrime. In other words, the rate of crime is not as bad as the picture often portrayed in the media and is, in some cases at least, even improving (Jardine 2015). In other words, law enforcement still has a reasonable chance, and is doing a fairly good job, of holding web-based crime at bay. Policing the dark Web can be successful.

CONCLUSION

Overall, Internet policing is maybe not ideal. It would be better if people just stopped using anonymity networks such as Tor to do illegal things. That would allow the network to be used to circumvent censorship and surveillance in repressive countries without any of the socially damaging spillover that online anonymity produces.

The network is fragile, despite its resilience, and if we try to find a quick and easy technological fix to problems that are actually social, we run the very real risk of breaking the Internet. Rather than discarding Tor or breaking the anonymity and encryption of the system through back doors for law enforcement, the focus should instead be on policing what goes on upon the network itself. Policing has the advantage of minimizing the costs that the dark Web imposes on society, while allowing the dark Web to have the maximum potential positive effect globally. It is not perfect, but it is the best we can probably do.

Acknowledgements

This chapter has benefited from the eyes of, in no particular order, Secretary Michael Chertoff, Laura DeNardis, Bill Graham, Gordon Smith, Pindar Wong and Leanna Ireland. It has also been polished by Carol Bonnett and Vivian Moser. Despite the collective wisdom of all these eyes, problems might linger. These are my errors alone.

WORKS CITED

- Abelson, Harold, Ross Anderson, Steven M. Bellovin, Josh Benaloh, Matt Blaze, Whitfield Diffie, John Gilmore, Matthew Green, Susan Landau, Peter G. Neumann, Ronald L. Rivest, Jeffrey I. Schiller, Bruce Schneier, Michael Specter and Daniel J. Weitzner. 2015. *Keys Under the Doormats: Mandating Insecurity by Requiring Government Access to All Data and Communications*. Computer Science and Artificial Intelligence Laboratory Technical Report. <http://dspace.mit.edu/bitstream/handle/1721.1/97690/MIT-CSAIL-TR-2015-026.pdf?sequence=8>.
- Bartlett, Jamie. 2014. *The Dark Net: Inside the Digital Underworld*. London, UK: William Heinemann.
- BBC. 2014a. "Child abuse image investigation leads to 660 arrests." BBC News, July 16. www.bbc.com/news/uk-28326128.
- . 2014b. "Russia Offers \$110,000 to Crack Tor anonymous Network." BBC News, July 28. www.bbc.com/news/technology-28526021.
- . 2015. "50 arrests in NI online abuse images probe in past year, say police." BBC News, March 15. www.bbc.com/news/uk-northern-ireland-31896685.
- Bradshaw, Samantha and Laura DeNardis. n.d. "The Politicization of the Internet's Domain Name System: Implications for Internet Security, Universality, and Freedom." Unpublished manuscript.
- Bradshaw, Samantha, Laura DeNardis, Fen Hampson, Eric Jardine and Mark Raymond. 2015. "The Emergence of Contention in Global Internet Governance." Global Commission on Internet Governance Paper Series No. 17. <https://ourinternet-files.s3.amazonaws.com/publications/no17.pdf>.
- Brandom, Russell. 2015. "Feds found Silk Road 2 servers after a six-month attack on Tor." The Verge, January 21. www.theverge.com/2015/1/21/7867471/fbi-found-silk-road-2-tor-anonymity-hack.
- Butterly, Amelia. 2013. "'Growing trend' of cyberbullying on social networks." BBC News, October 2. www.bbc.co.uk/newsbeat/article/24364361/growing-trend-of-cyberbullying-on-social-networks.
- Chertoff, Michael and Toby Simon. 2015. "The Impact of the Dark Web on Internet Governance and Cyber Security." Global Commission on Internet Governance Paper Series No. 6. https://ourinternet-files.s3.amazonaws.com/publications/GCIG_Paper_No6.pdf.

- Christin, Nicolas. 2012. "Traveling the Silk Road: A Measurement Analysis of a Large Anonymous Online Marketplace." Working paper, November 30. <http://arxiv.org/pdf/1207.7139.pdf>.
- Comey, James. C. 2014. "Speeches." www.fbi.gov/news/speeches/going-dark-are-technology-privacy-and-public-safety-on-a-collision-course.
- Cook, James. 2014. "FBI Arrests Former SpaceX Employee, Alleging He Ran The 'Deep Web' Drug Marketplace Silk Road 2.0." *Business Insider*, November 6. www.businessinsider.com/fbi-silk-road-seized-arrests-2014-11.
- Electronic Frontiers Foundation. n.d. "Anonymity." www.eff.org/issues/anonymity.
- Europol. 2011. "Operation Rescue." www.europol.europa.eu/content/operation-rescue.
- Freedom House. 2015. "Freedom in the World: Aggregate and Subcategory Scores." https://freedomhouse.org/report/freedom-world-aggregate-and-subcategory-scores#.Va6gr_IVhBc.
- Global Voices. n.d. "Anonymous Blogging with WordPress & Tor – ARCHIVED." *Global Voices*.
- Greenburg, Andy. 2014a. "The FBI Finally Says How It 'Legally' Pinpointed Silk Road's Server." *Wired*, September 5. www.wired.com/2014/09/the-fbi-finally-says-how-it-legally-pinpointed-silk-roads-server/.
- . 2014b. "FBI's Story of Finding Silk Road's Server Sounds a Lot Like Hacking." *Wired*, September 8. www.wired.com/2014/09/fbi-silk-road-hacking-question/.
- Human Rights Watch. 2006. *Race to the Bottom: Corporate Complicity in Chinese Internet Censorship*. Human Rights Watch. www.hrw.org/reports/2006/china0806/china0806webwcover.pdf.
- Jardine, Eric. 2015. "Global Cyberspace Is Safer than You Think: Real Trends in Cyberspace." Global Commission on Internet Governance Paper Series No. 16. https://ourinternet-files.s3.amazonaws.com/publications/no-16_Web.pdf.
- . n.d. "Tor, What is it Good For? Political Rights and the Use of Anonymity-Granting Technologies." Unpublished paper.
- Knibbs, Kate. 2014. "Silk Road 3 Is Already Up, But It's Not the Future of Darknet Drugs." *Gizmodo*, November 7. <http://gizmodo.com/silk-road-3-is-already-up-but-its-not-the-future-of-da-1655512490>.
- Koebler, Jason. 2012. "Online Black Market Drug Haven Sees Growth Double." *U.S. News*, August 7. www.usnews.com/news/articles/2012/08/07/online-black-market-drug-haven-sees-growth-double.
- Microsoft News Center. 2013. "Microsoft, the FBI, Europol and Industry Partners Disrupt the Notorious ZerAccess Botnet." December 5. <https://news.microsoft.com/2013/12/05/microsoft-the-fbi-europol-and-industry-partners-disrupt-the-notorious-zeroaccess-botnet/>.
- MIT Technology Review. 2012. "How China Blocks the Tor Anonymity Network." April 4. www.technologyreview.com/view/427413/how-china-blocks-the-tor-anonymity-network/.
- Murray, Andrew. 2014. "The dark web is not just for paedophiles, drug dealers and terrorists." *The Independent*, August 18. www.independent.co.uk/voices/comment/the-dark-web-is-not-just-for-paedophiles-drug-dealers-and-terrorists-9920667.html.
- NBC News. n.d. "Massive Online Pedophile Ring Busted by Cops." www.nbcnews.com/id/42108748/ns/us_news-crime_and_courts/t/massive-online-pedophile-ring-busted-cops/#.VcirBPIVhBc.
- Omand, David. Forthcoming. "The Dark Net: Policing the Internet's Underworld." *World Policy Journal* (Winter 2015/16).
- Owen, Gareth and Nick Savage. 2015. *The Tor Dark Net*. Global Commission for Internet Governance Paper Series No. 20.
- RCMP. 2014. "Project DISTRESS." www.rcmp-grc.gc.ca/mb/news-nouvelles/2014/project-projet-distress-20141211-eng.htm.
- Sullivan, Eileen. 2013. "Local Police Get into Cybercrime Fighting Business." *Huffington Post Tech*, April 13. www.huffingtonpost.com/2013/04/13/police-cybercrime_n_3075427.html.
- Thielman, Sam. 2015. "Silk Road operator Ross Ulbricht sentenced to life in prison." *The Guardian*, May 29. www.theguardian.com/technology/2015/may/29/silk-road-ross-ulbricht-sentenced.
- Woods, Andrew. 2015. *Data Beyond Borders: Mutual Legal Assistance in the Internet Age*. Global Network Initiative. January. <https://globalnetworkinitiative.org/sites/default/files/GNI%20MLAT%20Report.pdf>.
- Zetter, Kim. 2013. "How the Feds Took Down the Silk Road Drug Wonderland." *Wired*, November 18. www.wired.com/2013/11/silk-road/.

ABOUT THE AUTHOR

Eric Jardine joined CIGI as a research fellow in May 2014 in the Global Security & Politics Program. He contributes to CIGI's work on Internet governance, including the CIGI-Chatham House-sponsored Global Commission on Internet Governance. His current research focuses on cyber security, cyberterrorism, cybercrime and cyber protest. He holds a Ph.D. in international relations from the Norman Paterson School of International Affairs at Carleton University, Ottawa, Canada.

**CHAPTER FOUR:
THE TOR DARK NET**
Gareth Owen and Nick Savage

Copyright © 2015 by Gareth Owen and Nick Savage

BACKGROUND

The term “dark Net” is loosely defined, but most frequently refers to an area of the Internet only accessible by using an encryption tool called The Onion Router (Tor). Tor is a tool aimed at those desiring privacy online, although it frequently attracts those with criminal intentions. An innovative feature of Tor is the ability to host websites anonymously and with a degree of impunity — designed to be used by those in repressive regimes who wish to host whistle-blowing or political content.

The study described in this chapter collected data on the Tor dark Net over a period of six months to analyze the type and popularity of the content. Perhaps unsurprisingly, the majority of sites were criminally oriented, with drug marketplaces featuring prominently. Notably, however, it was found that sites hosting child abuse imagery were the most frequently requested.

INTRODUCTION

Tor is an open-source tool that aims to provide anonymity and privacy to those using the Internet. It prevents someone who is observing the user from identifying which sites they are visiting and it prevents the sites from identifying the user. Some users value Tor’s anonymity because it makes it difficult for governments to censor sites or content that may be hosted elsewhere in the world.

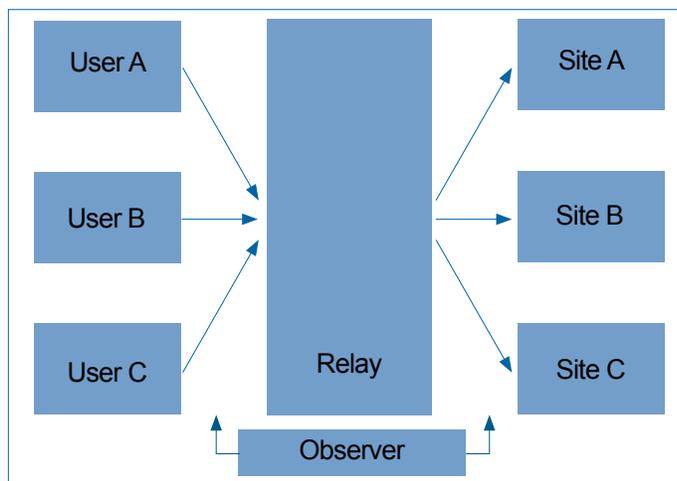
Tor has a critical mass of users, averaging two million per day as of June 2015 (Tor Project 2015), and is thus frequently cited as one of the key tools against government surveillance. Somewhat paradoxically, the Tor Project (the non-profit organization that manages Tor) receives the majority of its funding from the US government.

Tor volunteers run thousands of “relays,” a server that any other user can ask to route traffic through. Figure 1 illustrates the simple case of a single relay, with three users asking it to route traffic to three sites. An observer can see traffic entering and leaving the relay, but they cannot determine which user is visiting which site (save for correlation attacks, which will be discussed later) because the traffic is encrypted; however, if the relay operator is malicious, they can trivially (with ease, from a technical standpoint) link the two.

When a user visits any sites through a relay, his traffic appears to come from the relay rather than the user’s computer. Thus, the user remains anonymous to the site itself.

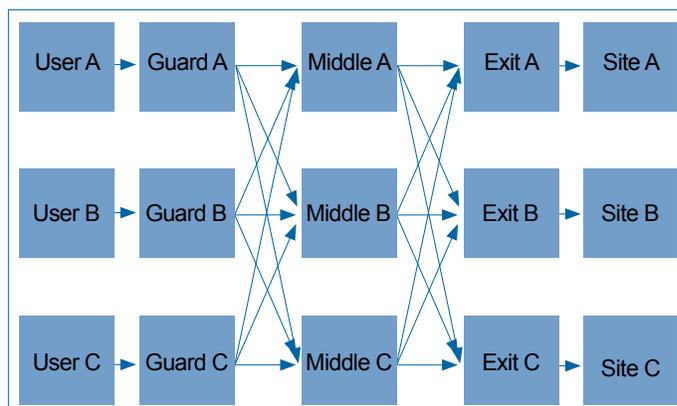
To defend against a malicious relay operator, the user chooses three relays and chains them together, labelling them as the Guard, Middle and Exit, known as a three-hop circuit (see Figure 2). This raises the bar significantly for an attacker who would then need to control all three

Figure 1: Illustration of Relay-mixing Traffic



Source: Authors.

Figure 2: Typical Three-hop Tor Circuits



Source: Authors.

relays to be able to link users with certainty to the sites they are visiting. It is the user who chooses the three relays; the attacker is unable to influence his choice/decision. An attacker’s only option would thus be to control a significant number of relays in the hope that a user chooses three within that controlled pool — this is thought to be impractical.

HIDDEN SERVICES

While the ability to access the Internet anonymously is valuable in countries where personal freedoms are restricted, it is only one feature of Tor (Jardine, n.d.). The other major feature is “hidden services” (HSes), the ability to host a website (or Internet service) anonymously. In this case, both the visitor and the site are anonymous to each other. Using this feature, political blogs or fora can be hosted in repressive regimes without fear of penalty. As with any technology such as this, it also allows the possibility of criminally oriented material to be hosted

with a degree of impunity. The collection of Tor HSEs is often referred to as the dark Net, although there exist other, less popular tools that might also be considered under this umbrella (for example, the Invisible Internet Project, known as I2P).

It is crucial to know how hidden services work to be able to understand the methodology used in measuring activity on the dark Net. Let's assume Bob is hosting an HS and Alice wishes to visit his site. When Bob first creates his site, he constructs a document detailing "introduction points," or relays within the network that will be able to relay messages to him. He publishes this document in a distributed hash table (DHT), which can be thought of as a database or phone directory distributed across all the relays in the network — that is, no single relay controls or possesses all of the DHT at any one point in time. To create such a database, all the relays in the network are placed onto a circle and ordered according to a unique identifier (see Figure 3, relays are labelled h_i). The HS is then mapped onto the circle at two points. Bob publishes information on his introduction points to the three relays to the right of each of these two locations, so that copies exist on exactly six relays. When publishing the information, he uses a three-hop circuit to remain anonymous to the directory relay.

The location that Bob publishes to in this directory appears random, and changes every day, but it is possible for Alice to figure out his information. The location changes daily to make it more difficult for one person to control the relays that hold Bob's information.

Alice then calculates which relays on the circle contain Bob's information and builds three-hop circuits to the relays, requesting a copy of his information. By using a three-hop circuit, she remains anonymous to the directory relays. She now has information on Bob's introduction

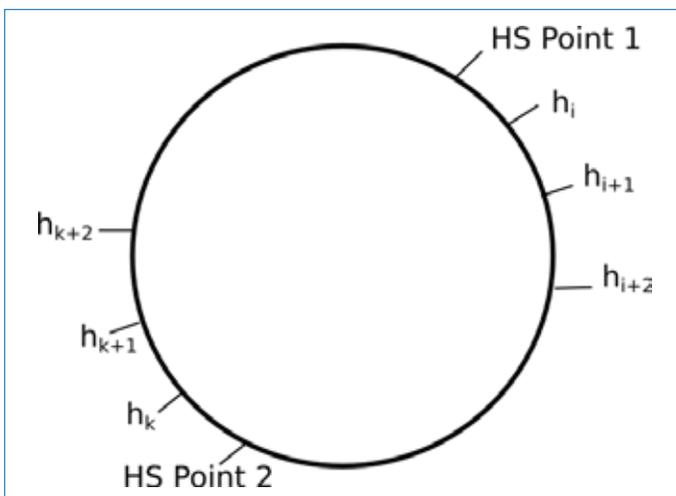
points and relays a message to one of them asking Bob to build a connection to a rendezvous relay that she chooses. The rendezvous relay proceeds to relay messages between Bob and Alice, and since both have connected to the rendezvous relay through three-hop circuits, the relay does not know the identity of either party. The rendezvous relay cannot inspect the traffic because it is encrypted; its service is wholly altruistic.

RELATED WORK

HSEs were described in the original Tor Paper (Dingledine, Mathewson and Syverson 2004) and have since undergone several revisions. They are difficult to locate geographically, but they use a DHT (similar to those used in many other distributed Internet applications (Stoica et al. 2001) to publish descriptors with information on how to connect to them. The Tor DHT is not resistant to Sybil attacks (Douceur 2002), in which one can run many nodes and gain control of a large proportion of the DHT. With that control, one can collect HS descriptors (as described in this study) and deny service to legitimate users. There exist a number of Sybil-resistant DHT implementations (Lesniewski-Laas and Kaashoek 2010), but as of yet, Tor has not focused significantly on this aspect.

The Tor DHT consists of approximately 3,000 participating relays, and each must have been operating persistently for several days before it can participate in the DHT. Furthermore, one can operate only two relays per Internet Protocol (IP) address to increase the cost of launching a Sybil attack. Researchers at Luxembourg University (see Alex Biryukov, Ivan Pustogarov and Ralf-Philipp Weinmann 2013) describe a bug in the Tor core program that allows someone to launch a large number of relays on a single computer and selectively phase any into the network. Tor logs relays' uptime, even if they were not in the network, thus making it possible to launch a number of relays on a single computer. Biryukov, Pustogarov and Weinmann took advantage of this bug and were able to collect the list of HS addresses in fewer than two days; however, the Tor Project has now fixed this bug. The authors used an automated classification algorithm to classify hidden sites into categories by content type. Their data shows they encountered popular abuse sites but chose to label them as "Adult" — an unfortunate side effect of the classification technique used. Additionally, they only examined HSEs present during a single 24-hour period. Therefore, the general question of the size, content and popularity of the dark Net remains open. This chapter addresses the question by collecting data over a significantly longer period of time and manually classifying sites to achieve greater precision.

Figure 3: Tor HS Directory (DHT)



Source: Authors.

STUDY OF HSEs

To collect information on the dark Net, a list of HSEs must first be enumerated. By controlling all the relays in the Tor DHT, it is possible to collect a complete list of HSEs by recording the descriptors as they have been published. It is then possible to count the number of requests for each descriptor and estimate their relative popularity.

Unfortunately, there are approximately 3,000 relays in the DHT, and even though one can create relays that participate in the DHT, it is impossible to control all of it. There is also a non-negligible cost associated with each participating relay one wishes to run. If one runs a handful of DHT participants, then one can observe a fraction of it at any one point in turn. Bearing in mind that HSEs publish to two essentially random points every day, over time one would observe every HS that remained online during the collection period.

In this study, 40 relays were operated for a period of six months. Each relay recorded a list of published HS descriptors and the number of requests for each. Although only a small proportion of the DHT was observed each day, cumulatively all of the DHT was observed many times throughout the study.

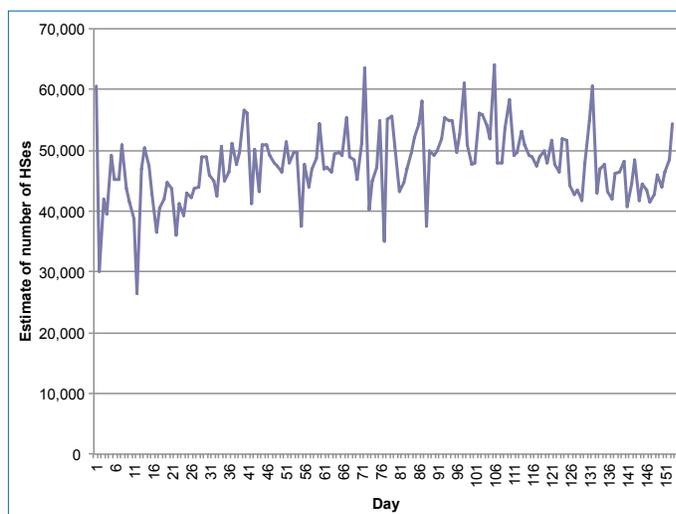
Size and Turnover Estimation

Little has been known about the dark Net to date, although a 2013 study estimated that there were 60,000 HSEs at any one time (ibid.); however, this study was based on a single day and it was not known whether this was an outlier. Extrapolating from this paper's data for each day, it is possible to give an estimate for the number of HSEs existing on any one day (see Figure 4).

While on first observation it looks as if the number of HSEs has high variance from day to day, one must bear in mind that only a small proportion of the DHT is being observed and then extrapolated. This means that errors will be amplified and this accounts for the variance. The long-term average throughout the study was 45,000 active sites and this is likely to be more indicative of the total number of HSEs. In total, 80,000 unique HSEs were observed during the study, but some only existed for a short period of time.

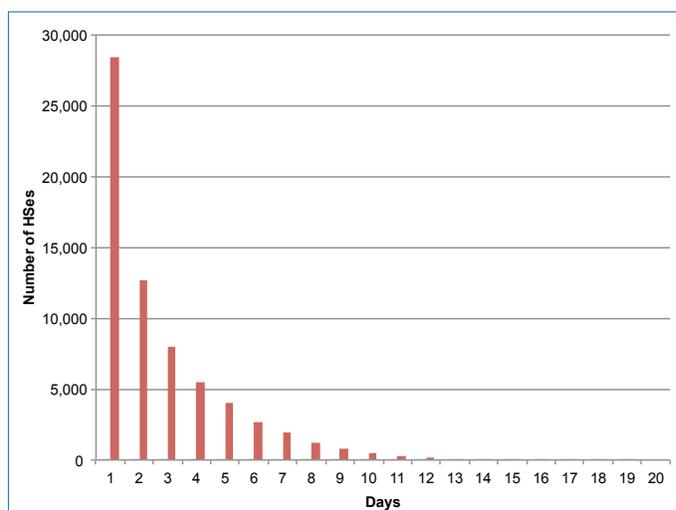
While observing a fraction of the circle throughout the study, and bearing in mind that an HS publishes to two random points on the DHT circle each day, one would expect to see a long-lived HS publish again and again to the relays. Figure 5 shows the number of days an HS was observed during the study. The largest number of HSEs were only seen once, which suggests that they existed for a short period of time and were never seen again. Longer-lived HSEs accounted for only 15 percent of all HSEs. Therefore, one can conclude that while there are many

Figure 4: Estimate of the Number of HSEs on Each Day of the Study



Source: Authors.

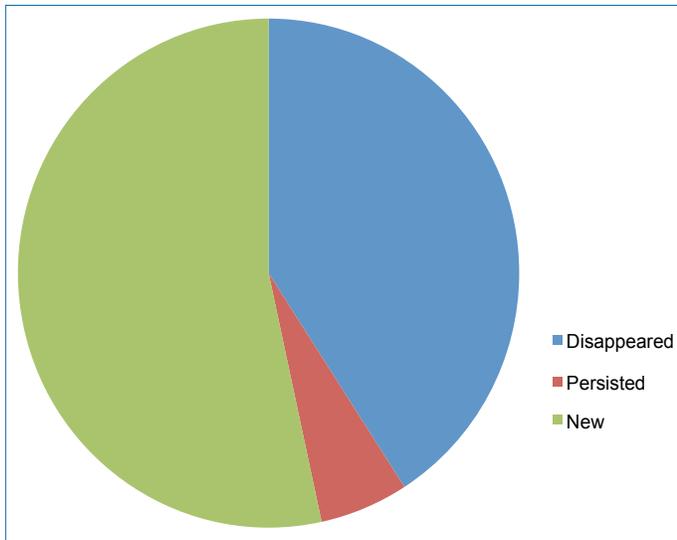
Figure 5: Number of Days an HS was Observed During the Study



Source: Authors.

HSEs, most only exist for a short period of time and are not long-lived services. The reason for this is unknown.

To further confirm the hypothesis that HSEs have a high turnover, the authors of an earlier study (Biryukov, Pustogarov and Weinmann 2013) were approached to provide their collected list. Figure 6 shows the number of HSEs that had disappeared, persisted or been created between the two studies. As one might expect, given the previous results, most HSEs did not persist for long and many newer ones had replaced them.

Figure 6: HSEs that Persisted More than 18 Months

Source: Authors.

Authenticated HSEs

Tor allows HSEs to be authenticated in such a way that one cannot locate the position on the DHT circle without knowing a secret, such as a password; therefore, unless one knows the secret, the service cannot be accessed. On the directory relays, it is trivial to identify the descriptors belonging to authenticated HSEs, because they are encrypted with the secret and so one can simply count encrypted descriptors (although not decrypt them). During this study, only 0.6 percent of HSEs were authenticated. The content of these HSEs is unknown, as without the secret it is impossible to access them.

CONTENT AND POPULARITY ANALYSIS

Classification

An HS does not have to be a website but could be, for example, a chat room, a file server or any other form of Internet service. There is no mechanism in Tor to find out which services are available for use by visitors, and the only way to discover them is to try each in turn (see Biryukov, Pustogarov and Weinmann 2013 for further analysis). Webpages were nearly universally offered by the most frequently requested HSEs. To identify the type of content available, a custom crawler was developed that would connect to each HS, download web content and extract key data points. These data points were then used for classification of content type.

Classification of webpages is a difficult task, and while there exist automatic classifiers based on machine learning, the dataset in this case was small enough that manual classification was not unduly onerous. Additionally, the

authors felt that given the range and complex, technical nature of some of the content, automatic classifiers would be insufficient due to difficulty in interpreting context and meaning (Samarawickrama and Jayaratne 2011).

Deciding when to crawl is not straightforward. At first glance, one may assume that crawling throughout the study is the logical approach; however, in doing so, one will overrepresent short-lived HSEs, most of which were not online concurrently. Instead, the preferred approach is to take a snapshot of the content at a particular period in time; having observed the turnover and short longevity of services, the crawling took place over a one-month period toward the end of the study. It is acknowledged that there is room for imprecision, but there is presently no better approach available. As there is a high turnover of HSEs, this is not significant.

It was considered at the outset that some HSEs may contain content that was obscene or otherwise illegal to download, and it was more than likely that if this content existed it would be in the form of multimedia or images; hence, the crawler only fetched textual content from each HS, parsed key data points and stored them in a database. The crawled data were inspected to produce a list of categories that covered the majority of the content. Afterward, each site was manually examined and classified into distinct categories. Where a site spanned two categories, the authors chose the category that more precisely described the overwhelming or primary purpose.

While there is often debate about the division of illegal and legal content on the dark Net, it is difficult to classify sites into either legal or illegal due to discrepancies and intricacies between legal jurisdictions: for example, whistle-blowing sites are often considered legal, but may not be if used to disclose classified documents or by persons in repressive regimes. Therefore, classification into legal versus illegal has not been undertaken. That said, the majority of sites on the Tor dark Net are likely to be illegal (or considered immoral) in many Western countries.

The classification categories are as follows, with notable examples where appropriate:

- **Abuse:** sites where the title indicates some form of sexual abuse (typically minors), likely to be illegal in most Western jurisdictions. Sadly, these pages were easily identifiable from the metadata, suggesting webmasters had confidence that Tor would provide robust anonymity. For some sites, it was difficult to discern whether they were facilitating abuse or providing adult pornographic services, and due to legal restrictions we were unwilling to download images to confirm. Where this was the case, the site was put into the porn category.

- **Anonymity:** sites aimed at promoting (or teaching) the use of anonymity tools or anonymous culture.
- **Bitcoin:** currency exchange from a mainstream currency to bitcoin, but more often money-laundering services.
- **Blog:** personal or topical blog, often covering topics such as hacktivism.
- **Books:** ebook service typically offering copyrighted material for free.
- **Chat:** web-based chat service, excluding services such as Jabber and Internet Relay Chat.
- **Counterfeit:** sites offering counterfeit items; notable fake currency, such as notes, or fake passports/identity documents.
- **Directory:** site offering links to other sites within the dark Net, often used for discovering other sites.
- **Drugs:** the sale or purchase of narcotics; typically, marketplaces connecting buyers and sellers.
- **Forum:** web-based forum whose primary purpose does not fit into another category; for example, generalist forum.
- **Fraud:** sites attempting to obtain a pecuniary advantage by deception.
- **Gambling:** any site that promotes/supports gambling. Bitcoin gambling services were most prevalent here, whereby users would first convert their fiat currency to bitcoin.
- **Guns:** sites exclusively aimed at selling guns.
- **Hacking:** site providing instructional information on illegal computer hacking.
- **Hosting:** dark Net hosting service allowing users to host another dark Net site.
- **Mail:** dark Net web-based email or messaging service; examples include Mail2Tor and the now defunct TorMail.
- **Market:** a marketplace selling items other than drugs or services covered in other categories.
- **News:** news service such as current affairs or news specific to the dark Net.
- **Porn:** pornography sites that carry material that would be legal in most Western jurisdictions.
- **Search:** site providing a search engine-type service; one example is Ahmia.

- **Whistle-blower:** sites typically operated by journalists for whistle-blowers to submit documents. The GlobaLeaks platform (Hermes Center for Transparency and Digital Human Rights 2014) and SecureDrop platform (Freedom of the Press Foundation 2014) were prominently featured in this category.

- **Wiki:** user-editable content, such as the Hidden Wiki.

Popularity

The popularity data here shows the number of requests made for the descriptor for a particular HS. Bearing in mind the earlier point about the Tor program caching descriptors, one can interpret the number of requests as between the number of visits and the number of visitors. Due to the anonymity offered by Tor, it is not possible to link two separate requests to the same person, but since their computer will remember descriptors until the Tor software is restarted, it often will not make multiple requests within a 24-hour period.

Table 1 shows the popularity of HSEs for which the authors received a descriptor request, but did not receive a publication during the study. These are addresses that no longer exist, but are still being requested by Tor clients. In many cases, it was possible to identify the purpose of these now extinct HSEs by examining online malware reports or by word prefixes present in the .onion address.

Almost all the top 40 HSEs requested but no longer operating were botnet command and control (C&C) servers (Stone-Gross et al. 2009). Botnet C&C servers are used to control computers infected with malware (called bots) remotely; the bot will connect to the server regularly for new instructions or to upload data (such as stolen passwords). Malware authors and researchers have been involved in a cat-and-mouse game in the last 10 years, whereby authors have attempted to produce C&C servers that are difficult to take down. Tor has become a popular tool for C&C infrastructure, due to the difficulty in taking down and locating servers. Interestingly, most botnets represented in the dataset had many (as opposed to a single) HS addresses, which paradoxically may make

Table 1: Popularity of No-longer Existent HSEs

HS Address	Requests/day	Days observed	Description
177ukkijdca2tsy	679,470	9	Botnet Sefnit
7sc6xyn3rrxtknu6	525,930	11	Botnet Sefnit
pomyeasfnmtn544p	514,766	10	Botnet Sefnit
ceif2rmdoput3wjh	247,296	6	Botnet Sefnit
censored	6,603	10	Child abuse

Source: Authors.

Table 2: Non-sequential Snapshot of Popular HSEs

HS Address	Requests/day	Days observed	Description
censored	168,152	12	Child abuse
silkroad6ownowfk	8,067	11	Silk Road
agorabasakxmewww	3,035	8	Agora
k5zq47j6wd3wdvjq	2,589	5	Evolution
xmh57jrznw6insl	1,341	7	Torch
3g2upl4pq6kufc4m	1,223	4	DuckDuckGo
wikitjerrta4ggz4	555	12	HiddenWiki
mail2tor2zyjdctd	266	8	Mail

Source: Authors.

them more vulnerable to deanonymization attacks if these services are distributed across several Tor processes (Murdoch and Zieliński 2007; Johnson et al. 2013).

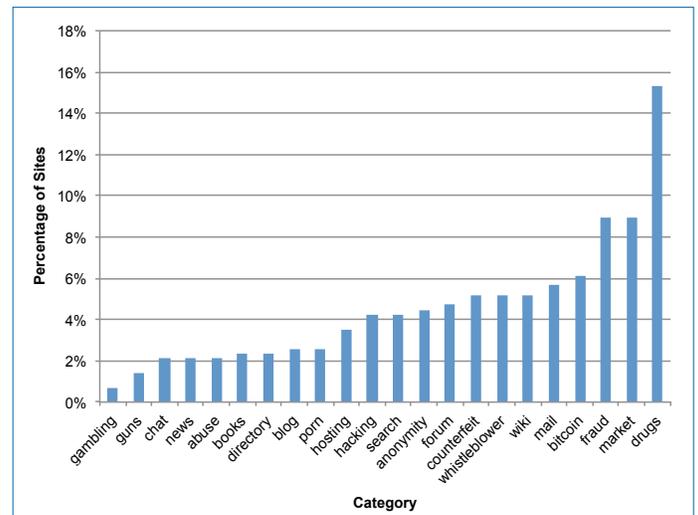
Table 2 shows a cross-section of the widely known .onion addresses by the number of visitors they received each day. Abuse sites were by far the most popular and these sites were easily identified by words in the page title or by prefixes used on the .onion address. The Hidden Wiki also featured and is often used as a starting point for many visitors to the dark Net. It is perhaps surprising, given the amount of media attention that Silk Road receives, that the number of its requests is fewer than 10,000.

Classification

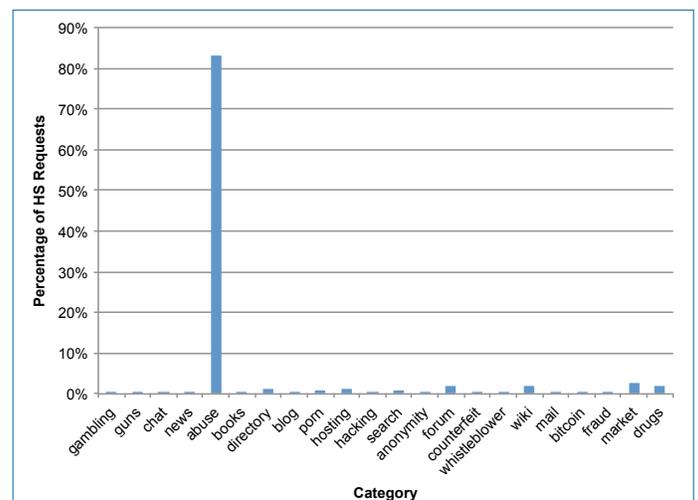
There are two representations of the classifications of data, the first being the number of sites in each category (see Figure 7). Figure 7 shows that the dark Net's content is diverse, with the largest number of sites being represented in the drugs category, but only by a small margin.

When each category is plotted against the percentage of HS directory requests it received (using the previous hits data), an entirely different picture emerges (see Figure 8). Requests to abuse sites represented more than 80 percent of total requests observed, although they accounted for only two percent of the total HSEs available (see Figure 7).

It is important to emphasize what is being measured. The popularity data is a measure of the number of HS directory requests, and when grouped into content-type categories the picture may become somewhat misleading. First, law enforcement frequently patrol abuse sites and this may inflate the figures; however, crawlers are likely to account for a single request in a 24-hour period and we are seeing a large number of requests to these sites. Even if it is assumed that all national forces crawl these sites daily, they would still only account for a small proportion of the total requests. The second possibility is denial of service attacks, where one could flood the HS directory with requests for descriptors in an effort to take the directory

Figure 7: Percentage of Sites in each Classification Category

Source: Authors.

Figure 8: Percentage of Requests by Classification Category

Source: Authors.

offline. This is likely to be ineffective because the attacker would need to take all six directories offline and then these relays would be dropped from the consensus and the responsibility would shift to other relays. It is worth noting that most of these sites were observed on several random days during the study, so an attack of this nature would have to persist for most of the duration of the study. While denial of service attacks are impossible to rule out, due to the anonymity offered by Tor, it seems unlikely and in any case none of our servers were taken offline or received requests far exceeding expectations.

Tor offers a tool called Tor2Web, which allows non-Tor users to visit HSEs through a web gateway. These web

gateways will operate one or a small number of Tor clients, so although there might be several visitors to a site, only one request will be seen because the gateway has cached the descriptor; hence, it is possible for some sites to be underrepresented in the data if they are largely accessed through Tor2Web. The popularity data is the proportion of HS directory requests observed on HSEs offering a website that the authors were able to crawl and classify. One should be extremely cautious before trying to link this data to a number of users, as the data approximates somewhere between visits and visitors. Interpreting the figure as visits will underestimate the number of users.

Connectivity

For each HS website that was crawled, information was extracted on the hyperlinks listed on their site. Each link was categorized into one of three categories: dark Net, clearnet or own-site. Dark Net links were links to other Tor HSEs, clearnet links were links to websites that were hosted on the Internet (for example, regular non-Tor domains) and own-site were links within the site being crawled.

Of the HSEs that were crawled, 59 percent did not link to any site other than itself, seven percent linked only to other dark Net sites, 23 percent linked only to clearnet sites and 11 percent linked to both. Aggregating the first two figures, one can say that two-thirds of sites were not connected by links to any sites outside of the dark Net.

DEANONYMIZATION OF TOR USERS AND HSEs

A common misconception is that Tor is resistant to state-level surveillance and that its users can therefore act with impunity. In reality, any suitably resourced entity can launch attacks with high success rates while maintaining a minimal risk of detection.

While an observer cannot see where traffic is routed in the Tor network, he can treat the network as a black box and observe traffic entering and leaving it. An interesting analogy would be the postal service, whereby one cannot see what happens in the sorting office, but can see how many letters/parcels every address posts and receives on each day. Assume the intelligence services think that two people are pen pals: they can observe letters leaving one person and arriving at the other and vice versa. Observing the mail of both parties over a period of time can give a degree of confidence about whether they are communicating with each other without opening their mail or tracking it through the postal system.

The postal system analogy may seem like there is a lot of room for error, but with Tor, a typical user may send millions of letters and an observer can see the precise time they were sent and received. It is therefore easy to confirm with high probability that two parties are communicating.

A slightly harder version of this problem occurs when one can observe traffic exiting the Tor network to a jihadist website: can he identify the original user while still treating Tor as a black box? The answer is often yes, provided he has enough visibility of traffic entering the Tor network to correlate the number of messages, the rate and time at which they are sent. One does not need to control guard relays to be able to launch traffic correlation attacks; one needs to be able to observe traffic between a user and his guard even though the traffic cannot be read at that point. Recent leaks from Edward Snowden indicate that UK and US intelligence services can observe traffic from entire countries, enabling them to observe all guards within those countries. Guards are presently changed every 30 to 90 days, so a targeted user may fall within the net at some point in the future when global observation is not possible.

While HSEs are believed to be the Holy Grail in anonymity protection, in reality these correlation attacks are much more successful compared with attacks against general web browsing through Tor. Typically with HSEs, one wishes to deanonymize the visitors and the service itself. In the last example, with general Tor usage, one was observing traffic entering and leaving the Tor network, while with HSEs the attacker can control one end of the connection and inject patterns of traffic to spot. In the case of a user, the attacker can control the relay in the DHT and send a specific pattern back to the user and try to identify it leaving the network. The attacker will be able to identify the service visited and the user but not what the user does on the site, because the content is encrypted end-to-end between the two parties. In some cases, the mere fact that a user has visited an HS may be enough to gain a conviction (for example, in particular where the site contains illegal content on the front page). That said, once a user has been identified, his home and equipment can be searched, where there may be stored evidence of wrongdoing. Thus far, the authors are aware of no cases whereby a deanonymization attack alone has been used to seek a conviction.

In the case of the service, the attacker can simply connect to the HS and send a pattern, and again attempt to identify it leaving the Tor network. Deanonymization attacks against HSEs can be highly successful with very low (even absent) false positives.

BLOCKING OF TOR

Tor is often described as being censorship resistant and impossible to block — this is not the case. There is a misunderstanding of how Tor works and some nations have attempted naive approaches that have, predictably, failed. There are many effective approaches to blocking Tor and the problem of building a truly censorship-resistant network is presently an open one.

When a user wishes to connect to the Tor network, he needs a list of all the relays, which he obtains from the

directory authorities — special relays operated by the Tor developers. He then selects a relay from the list that meets certain characteristics (principally uptime and bandwidth) and chooses this as the first node in any circuit. Since the list of relays (known as the consensus) is public, anyone is able to download the list and block access to all of them. The user would then be unable to connect to the first hop and into the network.

An attempt to mitigate these blocking attempts was made through the introduction of “bridges.” Bridges are not listed in the consensus and one has to visit the Tor Project website and enter a CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) to obtain a small number of them. The Tor website will only release a small fraction of bridges to any one user on any one day, which makes it difficult for an attacker to obtain the full list and block them.

Notably, China operates a country-wide firewall used to censor material that citizens can access. China has attempted to block Tor by stopping access to all of the relays listed in the consensus as described above (Winter and Lindskog 2012). They have also attempted to block bridges by looking for connections traversing the firewall to see if they met characteristics typical of Tor. This worked for some time until the Tor developers modified the program so that Tor connections were indistinguishable from ordinary web traffic. In response to this, China then monitored any encrypted connections and would try to connect to the remote server and talk to the Tor protocol: if it responded, they concluded it was a Tor relay; if not, it was an ordinary website.

There exist, however, many more successful techniques for bridge enumeration (Ling, Luo and Yang 2012) — that is, detecting potential bridges without asking the Tor Project for a list or scanning suspect servers. A simple approach is to run a Tor relay and monitor all of the circuits built through your relay. It is easy to identify whether you are the middle relay, so you can simply identify the previous hop and if it is not in the network it is probably a bridge. This technique is not foolproof because not all bridges will connect through your relay all of the time; hence, you must run many relays offering a significant proportion of the bandwidth to detect most of them, and even then you will only detect most, but not all.

That said, it is believed that these techniques would be effective, as only the most determined user would continue to persist with a tool that failed most of the time.

HS BLOCKING

While Tor is designed to be resistant to censorship, at present HSEs are not particularly robust against technical attacks (they will resist physical attacks if the operator is unknown). At present, groups of individuals or the

Tor Project itself could choose to block these sites by the following methods:

- An individual can block a single site by launching several relays and ensuring they occupy the positions in the DHT of the responsible relays for that service. If someone comes to the relay asking for the descriptor, the individual can simply deny it.
- Operators of Tor relays could themselves choose to block the content by patching their relays to deny requests to these sites. This would require the cooperation of a large percentage of relay operators to be effective, but it would be a decentralized blocking mechanism requiring some consensus.

The Tor Project itself can choose to trivially block the content by modifying the Tor program to block requests for such sites at the relay and client level. This might seem to place a large amount of power into the developer’s hands, but it is worth remembering they already control the authorities and the consensus, and can abuse this to deanonymize users or block sites anyway. At present, the Tor Project has stated that it is not willing to censor HSEs, because it fears it will be a slippery slope with future requests widening the categories blocked. This is unfortunate because child abuse sites do cause real harm and may encourage offenders. The number of requests for whistle-blowing sites is minuscule in comparison to those aimed at child abusers.

CONCLUSION

This chapter does suggest that child abuse content is the most popular type of content on the Tor dark Net. While law enforcement may crawl such sites, the number of requests that would be seen would be only a tiny fraction, and hence not skew the outline ratios. Similarly, denial of service attacks were not observed and so are also unlikely to account for the high requests. The usage of Tor2Web may underrepresent some categories, but it is not currently clear whether, or why, such groups would exclusively use this tool.

An explicit categorization of sites into illegal and legal was not undertaken, but it was abundantly clear to the authors that the majority of sites were of questionable legality. While anonymity and privacy tools such as Tor might fight online surveillance, they also give an easy and accessible route for those with criminal motivations. There are alternatives, such as botnets, available for criminal activity, but these do not negate the comparative ease with which Tor can be used.

It is technically possible to block Tor, although it is likely that the Tor Project will deploy countermeasures resulting in the endeavour descending into a cat-and-mouse game of “circumvent-and-censor.” In any case, Tor does not provide the absolute impunity that is often attributed to it.

WORKS CITED

- Biryukov, Alex., I. Ivan Pustogarov and Ralf-Philipp Weinmann. 2013. "Trawling for Tor Hidden Services: Detection, Measurement and De-anonymisation." *Proceedings of IEEE Symposium on Security and Privacy*: 80–94.
- Dingledine, Roger, Nick Mathewson and Paul Syverson. 2004. "Tor: The Second-Generation Onion Router." Washington, DC: Naval Research Laboratory.
- Douceur, John. R. 2002. "The Sybil Attack." *Revised Papers from the First International Workshop on Peer-to-Peer Systems*: 251–60.
- Freedom of the Press Foundation. 2014. "SecureDrop Platform."
- Hermes Center for Transparency and Digital Human Rights. 2014. "GlobaLeaks Platform."
- Jardine, Eric. n.d. "Tor, What Is It Good For? Political Rights and Online Anonymity-Granting Technologies." Unpublished manuscript.
- Johnson, Aaron, Chris Wacek, Rob Jansen, Michah Sherr and Paul Syverson. 2013. "Users Get Routed: Traffic Correlation on Tor by Realistic Adversaries." *Proceedings of the 20th ACM conference on Computer and Communications Security*.
- Lesniewski-Laas, Chris and M. Frans Kaashoek. 2010. "Whānau: A Sybil-proof Distributed Hash Table." *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*.
- Ling, Zhen, Junzhou Luo and Ming Yang. 2012. "Extensive Analysis and Large-scale Empirical Evaluation of Tor Bridge Discovery." IEEE INFOCOM. Orlando, FL.
- Murdoch, Steven J. and Piotr Zieliński. 2007. "Sampled Traffic Analysis by Internet-Exchange-Level Adversaries." *Proceedings of the Seventh Workshop on Privacy Enhancing Technologies*.
- Samarawickrama, Sameendra and Lackshman Jayaratne. 2011. "Automatic text classification and focused crawling." Sixth International Conference on Digital Information Management.
- Stoica, Ion, Robert Morris, David Karger, M. Frans Kaashoek and Hari Balakrishnan. 2001. "Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications." *Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*: 149–60.
- Stone-Gross, Brett, Marco Cova, Lorenzo Cavallaro, Bob Gilbert, Martin Szdlowski, Richard Kemmerer, Christopher Kruegel and Giovanni Vigna. 2009. "Your Botnet is My Botnet: Analysis of a Botnet Takeover." *Proceedings of the 16th ACM Conference on Computer and Communications Security*: 635–47.
- Tor Project. 2015. "Metrics Portal." June.
- Winter, Philipp and Stefan Lindskog. 2012. "How the Great Firewall of China is blocking Tor." *Proceedings of USENIX Workshop on Free and Open Communications on the Internet*.

ABOUT THE AUTHORS

Gareth Owen is a senior lecturer in the School of Computing at the University of Portsmouth. He holds a Ph.D. in computer science and has expertise in distributed computing systems, digital forensics and privacy-enhancing technologies. Before joining the university, he lectured at the universities of Kent and Greenwich in the United Kingdom.

Nick Savage is the head of the School of Computing at the University of Portsmouth. He was previously a principal lecturer in the School of Engineering at the University of Portsmouth, where he taught networking and security. He is a member of Working Group 3 for the European Commission's Network and Information Security Platform and has previously worked on projects funded by the Office of Communications and the Engineering and Physical Research Council. Nick holds a Ph.D. in telecommunications from the University of Portsmouth.

**CHAPTER FIVE:
CONNECTED CHOICES: HOW THE INTERNET
IS CHALLENGING SOVEREIGN DECISIONS**

Melissa E. Hathaway

Copyright © 2015 by Melissa E. Hathaway

ACRONYMS

BGP	Border Gateway Protocol
DNS	Domain Name System
ECJ	European Court of Justice
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IETF	Internet Engineering Task Force
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISO	International Organization for Standardization
ISPs	Internet service providers
ITRs	International Telecommunication Regulations
ITU	International Telecommunication Union
IXP	Internet Exchange Point
Mbps	megabits per second
OTT	over-the-top
Tor	The Onion Router
WCIT	World Conference on International Telecommunications

INTRODUCTION

Modern societies are in the middle of a strategic, multi-dimensional competition for money, power and control over all aspects of the Internet and the Internet economy. These struggles are occurring across a range of interrelated economic, technical, regulatory, political and social spheres, and the gamesmanship is intense. The players include multinational corporations, self-organized citizen and interest groups, and state and non-state actors. As such, these areas of tension are multilateral, multi-stakeholder and multicultural.

This competition has been increasing in focus, force and global reach since the birth of the Internet as an e-platform for commerce, information flows and power projection. In 1985, the potential for national power and wealth changed with the introduction of new top-level domains (such as .com). The Internet's potential became obvious in the early

1990s with the invention of the World Wide Web and was confirmed with broadband investments in the Internet's backbone network in the latter half of the decade. Today, the Internet community is able to click-connect-search-and-share information globally and almost instantaneously. The Internet facilitates access to and delivery of a wide range of services electronically, including e-government, e-banking, e-health, e-learning, next-generation power grids and air traffic control. The Internet also facilitates access to all things tangible, including military-grade weapons. The devices that connect people, places and things could offer up to US\$19 trillion in economic potential (Bilbao-Osorio, Dutta and Lanvin 2014); the modernization of industrial infrastructures already represents nearly 46 percent of the global economy — more than US\$32 trillion (Evans and Annunziata 2012, 13). As an instrument of power projection and military capability, today's networked systems, in particular the Internet, challenge traditional ideas of security, stability and sovereignty.

This infrastructure-Internet entanglement is a strategic vulnerability for all connected societies. The positive impact of the Internet on countries, communities, businesses and citizens can only be sustained if the service is accessible, available, affordable, secure, interoperable, resilient and stable. This is why the Internet and its underlying value proposition has become a national security matter. Global leaders must wrestle with the fact that their Internet infrastructures and citizen-facing services are vulnerable to interference and that their economic dependence on the Internet will not permit them to abandon the adoption path they are on (Hathaway 2010). They are also trying to diffuse or take advantage of the growing perception by many around the world that the United States has too much “control” over the Internet. The widespread view is that since the Internet was created in the United States, its companies dominate the information communications technology marketplace and are generating tremendous wealth for the West. Hence, the United States is perceived to be acting in its own interests to the detriment of others.

This chapter discusses the increasing pace of discord and the competing interests that are unfolding in the current debate concerning the control and governance of the Internet and its underlying infrastructure. Some countries are more prepared and committed than others to winning tactical battles on the road to becoming an Internet power. Some are acutely aware of what is at stake; the question is whether they will be the master or the victim of these multi-layered power struggles as subtle and not-so-subtle connected choices are made. Understanding this debate requires an understanding of the entangled economic, technical, regulatory, political and social interests implicated by the Internet. Those states that are prepared for and understand its multi-faceted nature will likely end up on top.

ECONOMIC INTERESTS

The importance of money flows from its being a link between the present and the future.

– Keynes (1935)

The first strategic area of competition is economic and concerns connectivity and infrastructure development. By the end of 2014, the Internet will be accessible to approximately 40 percent of the global population — most of whom are located in Western and more developed countries. The demand curve and market growth potential for connectivity and Internet penetration for the foreseeable future is likely to come from Asia, Africa and South America — with these come potential power and influence for their populations (Internet World Stats 2014). However, the predicate to Internet access is the provisioning of the underlying infrastructure that can deliver affordable broadband Internet services to citizens. Governments and companies are racing to lay the foundations for universal access for citizens, while simultaneously tying access to their economic sustainability and development agendas. This economic activity is being closely tracked by the International Telecommunication Union (ITU), the Inter-Development Bank, the Organisation for Economic Co-operation and Development and the World Bank, all of which have been ranking countries on their telecommunication initiatives.

Advancing connectivity requires promoting network and broadband infrastructure expansion. These investments can be costly — and countries may not have the means to deliver high-quality, low-cost infrastructure to remote areas with smaller populations. In the days of the landline telephone system, revenue was incurred through an inter-carrier international settlement system that negotiated a price per call based on origination and termination. This collection system helped pay for telecommunication infrastructure improvements aimed at reaching more and more citizens. However, in today's Internet Protocol (IP) environment, the concept of a "call" has no direct counterpart. Internet service providers (ISPs) may pay transit fees based on capacity or use settlement-free peering, thus bypassing the payment scheme previously imposed by inter-carrier international agreements. Content providers that offer their services via the networks of infrastructure operators using an over-the-top (OTT) model pose further challenges to this model. OTT content and services providers include Google, Facebook, PayPal, Amazon, Skype and others. These OTT providers consume bandwidth through their delivery of volumes of information to users transiting the infrastructure — usually for free. Sometimes these services can degrade the quality of the infrastructure operators' own telecommunication services, including core services, because they are using more than their "fair share" of bandwidth. Infrastructure operators are thus forced to

make additional investments to ensure that they can provide their customers with the low-latency, high-quality experience that they demand 24 hours a day, seven days a week. To further complicate matters, the majority of the OTT companies are headquartered in the United States.

Both national leaders and infrastructure operators feel threatened by this complex ecosystem. The entities that can "control" information flow can also assert or extract economic and political leverage. The perceived or very real inequality of who monetizes access to the Internet on the one hand, and who benefits from that access on the other, remains part of the ongoing debate. First, countries are seeking mechanisms to pair market access with cost-recoverable investments to pay for the infrastructure modernization that the twenty-first-century digital society is demanding. Some leaders are looking to the regulatory environment and international treaty venues, such as those convened by the ITU, to assert power over ISPs and OTT providers. Second, the market liberalization of the past two decades may give way to the resurgence of state-run telecommunications companies that, acting as ISPs, would be the conduit for citizens to reach the Internet. This gives nations more "control" over private or quasi-private providers, allowing them to channel the proceeds into their own economy. Depending on the argument made, this could be perceived as a barrier to market access. For example, the German government recently made a decision to phase out the use of Verizon Communications services by 2015 and transition to Deutsche Telekom to provide communications services for German government agencies. The change was made because of concerns about network security and citizen privacy (Troianovski and Yardin 2014).

A related aspect of the economic competition that has emerged around the Internet involves the movement of data across borders. For example, the Transatlantic Trade and Investment Partnership and the Trans-Pacific Partnership are regionally based free trade agreements, both of which are seeking to increase economic growth. The parties to these agreements will have to enable the free flow of data across borders if they wish to facilitate commerce. Yet, some countries are seeking mechanisms to protect their data, declaring that there needs to be data sovereignty for national security purposes. Can the data assume the "flag" of the country in which it was created?¹ The controversy is particularly challenging in an era where data is stored in multiple centres and geographic locations to enable citizen access on demand. This raises two fundamental legal and political questions. First, does the data assume the citizenship of its creator or of the country in which it is stored? Second, what happens when the data is shared

¹ Some countries are debating the merits of keeping data contained inside the geographic boundaries of their home country. If the data leaves the geographic borders, then it must be marked or "flagged" accordingly.

or backed up across multiple data centres in multiple geographic locations? The intermediaries — i.e., those who enable cross-border digital trade — will inevitably have an impact on national economies. They could also assert control in terms of influencing who benefits and who pays, thus presenting potential security challenges. For example, some countries may want to impose a jurisdictional right to inspect all data communications, while others may demand that organizations use indigenous “preferred” service providers and store data locally, thus forcing data to fall under local laws and giving potential access to law enforcement and intelligence services.

At the same time, efforts to promote the development of Internet Exchange Point (IXP) facilities to enable the quick transit of data through IP interconnections have accelerated. As countries strive to connect citizens in remote geographic locations, they will need multiple IXPs to ensure low-latency delivery, while striving to ensure end-to-end quality of service. Meeting these demands will also require operators of IXP facilities to take measures to further the security, safety, continuity, sustainability and robustness of their infrastructure. As a result, the companies or countries that build these IXPs will have a great deal of power over network traffic and the content that transits through those pipes.

The actors that dominate market access to, and provision of, the Internet will have the opportunity to assert control over information flows as well. If this power struggle continues along its current trajectory, future Internet growth will be dominated by the East and the South, and a new set of governments and constituents will seek to assert their voice, leverage and market power to achieve their own economic, political, military and societal goals. The United States heretofore has been perceived as the dominant player — perhaps even the colonial power of the Internet — not least because it has been the main developer and provider of Internet technologies and services. It is also perceived as being the main financial benefactor of the Internet. Today, however, the United States and its innovation centres of excellence are struggling for access and influence and may soon face displacement as new market leaders emerge around the globe.

TECHNICAL INTERESTS

The supreme art of war is to subdue the enemy without fighting.

– Sun Tzu (1963, chapter 3)

The second strategic area of competition that has emerged around the Internet is technical, involving multilateral decision-making bodies and multi-stakeholder processes. Both sets of constituents are debating who is best suited to govern the technologic foundations of the Internet. It is estimated that in the next five years, the Internet population

will double and the number of connected devices will reach at least 50 billion (Evans 2011). The effects of the “Internet of Things” — the devices that connect people, processes, data and things — will place considerable demands on existing institutions and governance mechanisms, some of which have long-standing practices and natural leaders. Competition over Internet-related technical interests is being waged on five fronts: infrastructure, protocols, standards, security and content.

Infrastructure

The underlying infrastructure of the Internet is constantly changing. ISPs come in many forms and sizes and go by many names: the phone company, the cable company, the wireless company, the satellite company and others. In the future, the Internet may be provisioned by an unmanned aerial vehicle or high-altitude balloons to connect those in rural and remote areas who have no Internet access.² ISPs are increasingly measured by their speed of service (for example, upload and download times at megabits per second [Mbps]). The most technologically advanced cities in the world enjoy speeds of up to 100 Mbps and hope to advance beyond 1,000 Mbps (Rediff Business 2013). In 2014, about 25 major ISPs carried 80 percent of the world’s Internet traffic. By 2020, this number will likely change as new delivery technologies emerge (such as unmanned aerial vehicles and balloons). Of course, these new technologies will have to navigate international politics as international conventions, administered by the ITU, determine allocation and use of the radio spectrum. These technologies may also come under scrutiny for their need to loiter in sovereign airspace (Fitchard 2013). So when companies such as Google expand their market position to gain more control of the Internet backbone to deliver their services without intermediaries, they should not be surprised that they face opposition. These new technologies and projects also threaten to displace the traditional providers (such as China Unicom, Nippon, Telefonica, Telegraph, Telephone, Telstra, Verizon and Vodaphone) that, in turn, are putting pressure on their governments and multilateral organizations to intervene to protect their interests. In some cases, defending the interests of the traditional providers is also convenient for the country because it advantages indigenous companies and enables the government to assert control over those who are trying to evade regulation and payment schemas.

Protocols

In addition to competition for the delivery path of the Internet, competition around how data moves through the Internet has also emerged, adding further complexity to the management of the Internet. First is the Domain

² Google launched Project Loon to use a global network of high-altitude balloons to connect people in rural and remote areas who have no Internet access. It began as a pilot in New Zealand and is expanding into Africa and elsewhere. See www.google.com/loon/.

Name System (DNS). Think of this as the “telephone directory” for the Internet in the sense that “[d]omain names are human-friendly names that are translated into [IP] addresses, for example, *www.acme.com* is a domain name, and *216.27.178.28* is its IP address” (Hathaway and Savage 2012, 15). Second are the individual protocols that are assigned to devices. The Internet of the twentieth century was designed to accommodate approximately 4.3 billion addresses (Bradner and Mankin 1995), and was enabled through the Internet Protocol version 4 (IPv4). The Internet of the twenty-first century, however, demands a much richer supply of addresses to accommodate the Internet of Things uptake and field more than 50 billion devices. It also requires the adoption of the Internet Protocol version 6 (IPv6) protocol, which will open up 340 trillion, trillion (3.4×10^{38}) unique addresses.

The transition to IPv6 poses at least two challenges. First, the providers of the transport layer — those who deliver the Internet service — will need to ensure interoperability between IPv4 and IPv6 devices. A translation mechanism is needed to enable IPv6-only hosts to reach IPv4 services and to allow isolated IPv6 hosts and networks to reach each other over IPv4-only infrastructure (SixXS 2015). This will require ISPs to invest in the necessary technology to enable a seamless experience for their global users. Developing this mechanism is not an insignificant cost. The second challenge derives from the nature and perceived “nationality” of the entity that is in charge of the global coordination of the DNS root, IP addressing and other IP resources — the Internet Assigned Numbers Authority (IANA), a department within the Internet Corporation for Assigned Names and Numbers (ICANN). The IANA functions are coordinated with and funded by the US Department of Commerce’s National Telecommunications and Information Administration. This perceived influence of the United States over the timing and allocation of Internet addresses and how the telephone directory of the Internet moves data is problematic.

The fact that the United States (via ICANN) is seen as controlling the protocols of the Internet is, indeed, the reason why many international venues are debating the merits of multi-stakeholder administration versus multilateral governance. Some countries believe that moving some functions of the Internet into a more global United Nations-like forum would ensure fairer distribution of the Internet resources needed for their digital societies. Russia and China are certainly lead advocates for this approach. Other countries, too, echo this call for global governance and are advocating for the Internet Governance Forum to be transformed into a World Internet Council and become the steward of the Internet (Euractiv 2014). Some global leaders posit that this would be more representative of their countries’, corporate and citizens’ interests and make how and why decisions are made more transparent. To diffuse the growing distrust in United States’ involvement

in the IANA functions, in March 2014 the US government announced its intent to transition its role and asked ICANN to convene global stakeholders to develop a proposal for that transition plan (US Department of Commerce 2014). Of course, this may not quell the desire to move the administration and governance of Internet resources into a multilateral venue.

Standards

The Internet society of the twenty-first century demands an interoperable Internet and devices that connect to that modernized infrastructure designed to work on any ISP backbone using standard protocols. This is where standard-setting bodies emerge as a strategic leverage point to influence the design specifications of the next generation of Internet products and services. There are a number of standards organizations, but two principal organizations that affect the global marketplace in this area.³ The first is the Internet Engineering Task Force (IETF), which manages the process of creating Internet standards. During this process, a “specification undergoes a period of development and several iterations of review by the Internet community and revision based upon experience, [and then it] is adopted as a Standard by the appropriate body and is published” (IETF 2015). The second organization is the International Organization for Standardization (ISO), an international standard-setting body comprising representatives from various national standards organizations. Its technical process leads to “endorsed” international standards that are often the benchmark that global corporations must design for and deliver to. Of course, there are many other standard-setting bodies, but these two affect much of the global Internet device and service market. Therefore, whoever designs these standards, creating that interoperability for global opt-in and global uptake, will also have a dominant presence in the market. Corporate and government players alike are positioning themselves to influence the outcomes of these two organizations because their decisions will determine market share, market influence and, subsequently, market control.

Security

Surveillance, piracy, criminal activity, intellectual property theft and physical harm/destruction are on the rise, with the Internet enabling much of it. As a result, securing the Internet infrastructure and the data and services that transit through it has become of paramount importance, sparking global debate and discord. Views differ on what is to be secured, how to secure it and who should perform the duties. Some countries are turning to ISPs, which

³ Other international standard-setting bodies include: the Institute for Electrical and Electronics Engineers, the Organization for the Advancement of Structured Information Standards and the World Wide Web Consortium.

have unparalleled access to global networks, to provide upstream security for downstream devices. Initiatives of this type include blocking spam seen in transit, identifying compromised devices owned by customers, quarantining infected devices and blocking their access to the Internet, identifying and blocking sources of distributed denial of service attacks, and minimizing frequency and duration of network outages and route disruption. But this represents only one layer of the current amalgam of security actors.

Others are advocating for a system to ensure the security and management of the DNS root. A single root is needed to ensure global uniqueness regarding names (both administration and allocation). Multiple roots might fragment the Internet, causing latency and misrouting, and potentially degrading Internet interoperability. As noted, some countries believe that the United States, through ICANN, is unfairly administering the system and they are arguing for an alternative, more regional or local system of governance with multiple roots. Their arguments are further fuelled by newspaper headlines about the United States' monitoring and surveillance practices, as well as its potential manipulation of data encryption standards (Jackson 2013).

The regionalized Internet argument has other security undertones that may affect data routing and OTT providers. For example, France and Germany are considering a Schengen routing system⁴ for data in Europe (Deutsche Welle 2014). But this raises another question: is the intention of this proposal to better protect the privacy of their citizens or is it to control digital trade and cross-border data flows? In 2012, for instance, Iran announced that it would pursue a national intranet, block services from Google, Yahoo Inc. and Hotmail, and replace them with indigenous and government-led programs such as "Iran mail" and "Iran search engine" — in line with Iran's plan for a "clean Internet" (Hathaway and Klimburg 2012). The emergence of other similar national intranets with national (non-Western) services is occurring more frequently, especially in the shadow of media reports about the scope of United States surveillance and intelligence-gathering activities.

Measures designed to secure the traffic, and the related infrastructure, come in many forms. Some are pushing for a DNS SECurity (DNS security extensions), which would make it possible to validate the authority of a query and response and ascertain whether the signed data has been changed during transport. The latter would limit interception and surveillance mechanisms. Others argue that ISPs should have a process or framework for securing Border Gateway Protocol (BGP) announcements — i.e.,

how data moves from one ISP to another — that includes specific technical procedures and protocols to ensure that routes cannot be "hijacked," rerouted or brought offline. In April 2010, for example, BGP users received an alert regarding a prefix hijack by China's largest ISP — China Telecom. Internet traffic was rerouted for approximately 15 minutes as a result, affecting both Chinese and American Internet traffic. This event "underscores the vulnerability of the BGP routing infrastructure and reminds us that if intentional, the criminal could store, alter or just throw away the traffic" (Hathaway and Savage 2012). The fact that BGP is vulnerable to hijack, and that it has been done on a number of occasions, has led to many countries wanting to know where all of their traffic has been and where it will be routed.

Still others are arguing for different protective measures for facilities, infrastructures and even content. Protective measures date back to the original 1934 International Telegraph Convention, which gives the ability to stop messages that "may appear dangerous to the safety of the State or would be contrary to the laws of the country, public order, or decency" (ITU n.d., Article 34). In 1988, public use of the Internet was in its infancy, and the International Telecommunication Regulations (ITRs) compiled that year did not contain explicit provisions for securing the traffic and supporting infrastructure.⁵ They did, however, include a reference for member states and the operating agencies to avoid "technical harm" (Article 9). This "special provision" was added as a reaction or afterthought at a time when member states were faced with the release and propagation of the Morris Worm that affected 10 percent of the Internet's computers and disrupted Internet services for days.⁶ Today, such a provision may translate into arguments to allow states to interfere with communications whose purpose is, indeed, to hinder the internal affairs or undermine the sovereignty, national security, territorial integrity and public safety of other states.

Security arguments are being used to empower governments to advance their economic, political and

4 This builds on the Schengen Agreement of Europe. There is a proposal (initiated by France and Germany) to regionalize the routing of European information to keep it in Europe — thus establishing borders for the flow of Internet traffic.

5 The ITRs were signed by 178 countries and are a recognized global treaty. The purpose of the treaty was to facilitate global interconnection and interoperability of international telecommunications networks by establishing a regulatory framework to: govern traffic flows between telecommunication network operators; address international routing, charging, accounting and billing between operators; assure quality of international services; and encourage avoidance of harm to networks and services. The regulations are credited with providing for economic growth via the e-economy and development around the world by liberalizing telecommunications and creating interoperability among network providers.

6 Robert Morris, Jr., a graduate student in Computer Science at Cornell University, wrote an experimental, self-replicating, self-propagating program called a "worm," and injected it into the Internet on November 2, 1988. Morris was tried and convicted under the Computer Fraud and Abuse Act of 1986.

military interests in the operational implementation and architecture of the Internet. This weakens multi-stakeholder processes and venues and, at the same time, boosts market access, disrupts the power and control over Internet governance, and positions states for standards leadership.

Content

Technology innovation over the last 20 years has also led to big changes in data generation, consumption and analysis. The modern digital society — both people and devices — is generating a lot of data. Looking at the widespread use of tablets, cellphones, cameras, EZ-passes for the highway, cars, smart grid, etc., we see that we live in a world of near-ubiquitous data generation. This reality is coupled with the declining cost of collection and storage, and new capabilities for processing and correlating data (The White House 2014a). Moreover, it has led to the emergence of new power brokers — those intermediaries who buy, sell and correlate data about citizen and device interaction with and over the Internet. Data aggregators amass online and offline information about people, culling details from websites, social media, search engines, buying habits, travel patterns and even government databases (The Federal Trade Commission 2014). They use technology and statistical algorithms to combine multiple sources of data to make inferences about individuals, their interests and the devices they use. They uncover patterns of activities and profile and track individuals — all this has profound implications for government and society, especially in terms of surveillance and censorship. This capability is no longer the sole purview of government intelligence services. In fact, with the right tools, commercial companies such as Google, Baidu, Facebook, Tuenti, Badoo and Renren are just as capable and have access to troves of data (Sorav 2012; TechWatch 2014).

The content of the flow of data over and throughout the Internet is important because it has significant economic, political and social value. Those who can tap into that content, therefore, have power. If mined and leveraged properly, this data can help identify the next consumer market (such as where to place the next Walmart), help locate suspected terrorists or dismantle an organized crime syndicate. It can also open new venues to exchange ideas and create new subjects for censorship. In the United States, law enforcement officials (such as the Federal Bureau of Investigation) use social platforms such as Facebook and Twitter to garner tips about suspected terrorists (Sterling 2012). On the other hand, some countries use citizens' digital footprints to search for and suppress those who might pose a threat to a regime's stability. For example, in March 2014, Turkish Prime Minister Recep Tayyip Erdoğan instructed ISPs operating in Turkey, including TurkTelekom, to seal off access to social media sites such as YouTube and Twitter (Zmijewski 2014). This action was

taken in response to Turkish citizens having used social media to organize protests across the country against his government's policies. In February 2014, Russia passed a new censorship law demanding that ISPs block access to websites deemed to contain information promoting extremism and/or endangering public safety. As noted by one commentator, the wording of this law can be broadly interpreted to "forbid pretty much anything critical of the ruling government: political opposition, environmental activism, provocative political art, investigative journalism, nonviolent political protest" (Levine 2014).

Countermeasures are also being fielded to circumvent increased surveillance and censorship. For example, The Onion Router (Tor) is free software and an open network that enables communications (and content) to move around a distributed network of relays run by volunteers all around the world who are circumventing measures to block their communications. It allows people and groups to increase their privacy and security on the Internet and keep some anonymity. Originally developed for the US Navy for the primary purpose of protecting government communications, Tor is now widely used by dissidents, activists, journalists, law enforcement personnel and military constituents. Some governments facilitate the use of Tor to enable freedom of speech and to promote democratic values. Those governments, however, are often criticized for interfering in the sovereign business of other states — namely in their regime legitimacy and stability. Of course, many other countries are trying to block the use of Tor (or crack its code) for national security purposes (The Tor Project 2015a).

Increasingly, we are seeing national leaders interfering with the Internet on behalf of their own interests (The Tor Project 2015b), with tensions rising between states as a result. Global leaders and citizens in different parts of the world are demanding clarification on data ownership, privacy and transparency. In short, they want to know what is being done with their data and how it is being used. In addition, many democracies continue to push for Internet freedoms and have declared access to the Internet a human right. More autocratic or authoritarian regimes, however, increasingly view the Internet as a threat. Others, like the United States, in a subtler and more hypocritical way, demand that other countries refrain from censoring their citizens while simultaneously pursuing their own broad-based monitoring and surveillance programs. This, in turn, does not help instill confidence in the legitimacy of the United States for Internet leadership.

Competition to shape the technological foundations of the Internet is strong — not least because it can lead to greater power, control and monetization of the Internet and the Internet economy. Its future is being debated in a range of international venues and bodies, ranging from the ITU to the IETF to ICANN and the ISO. How its functions and features should be governed is also being discussed

by entities like the World Economic Forum in special meetings like NETmundial (NETmundial 2014a) — which took place in April 2014 in Brazil — and by commissions like the Global Commission on Internet Governance (CIGI 2015). It is in these venues that the future course of infrastructure, protocols, standards, security and control of content will be determined.

REGULATORY INTERESTS

To widen the market and to narrow the competition is always the interest of the dealers.

– *Smith (1909, 14)*

The third strategic area of competition is regulatory, which is focused on ensuring that the Internet remains accessible, affordable, secure, stable and interoperable for everybody. Market mechanisms are being used to assert leverage and control, and to change the balance of power, politics and wealth creation. Countries and companies are at odds in this field. The subtle struggle is focused on how to govern the growth of the Internet — namely what is in the best interests of society and government versus what is in the best interests of companies and their shareholders. The main challenge lies in the fact that the private sector designs, builds, operates, maintains and restores the very systems that process, transmit and operate the country's most important information and most vital infrastructures, while governments remain the ultimate guarantor of their citizens' safety and well-being.

It is thus the responsibility of governments to facilitate the market to meet the economic and national security interests of their citizens. Most of the time this encompasses the provisioning of citizen-essential services like water, electricity and telephone access. Now that the Internet affects these and other citizen-essential services, governments are evaluating whether the Internet is in need of some sort of market corrections. The challenge, however, is establishing what exactly should be governed. Is it the functional areas of infrastructure provisioning, DNS administration, the standards-setting processes and the security thereof? Or is it the actual facilities, devices, companies and market access that need to be governed? Each country is using different market levers, in the form of legislation and regulations, to assert control, manage risk, build security back into the infrastructure and maintain political stability. For example, the European Parliament has released a draft legislative directive, "Measures to Ensure Network and Information Security" (European Commission 2013). This directive, if passed, would legally bind member states to be compliant with specific criteria, adopt appropriate steps to manage security risks, and report serious incidents to their national competent authorities. The directive is targeted to the operators of critical infrastructures, such as energy, transport,

financial services and health care, and to key providers of information society services, such as e-commerce platforms and social networks. The United States has signalled a similar intent to regulate broad industry sectors in Presidential Executive Order 13636, "Improving Critical Infrastructure Cybersecurity" (The White House 2013).

Other countries are turning instead to international treaty mechanisms to affect the market as well as contain political and social unrest. For example, in 2012, the ITU convened a World Conference on International Telecommunications (WCIT) to update and revise the ITRs (ITU 1989). The ITRs define the general principles for the provision, operation and compensation of international telecommunications services. The WCIT represented a perfect venue for countries seeking to assert more control over many aspects of the Internet, including facilitating an accounting mechanism to compensate for the infrastructure improvements needed to carry the ever-growing Internet (voice, data and video) traffic and to initiate security requirements for key facilities and networks. At the conclusion of the meeting, some 89 nations signed a new treaty and approximately 55 did not. The United States led the dissenting block, which had advocated for either maintaining the status quo or no change at all to existing ITRs, and has been criticized for its position ever since. At the time of negotiation, Ambassador Terry Kramer, the United States' lead negotiator, stated, "[w]e are disappointed with revisions that expand the treaty scope to Internet-related matters and content. We believe these provisions reflect an attempt by governments to regulate the Internet and its content, potentially paving the way for abuse of power, censorship and repression" (quoted in Rash 2012).

As one might expect, the debate or intent to govern the growth and assert control over the Internet did not end with the WCIT meeting in Dubai in 2012. In fact, many policy issues have extended into the discussions of the World Telecommunication/Information and Communication Technology Policy Forum, the World Telecommunication Standardization Assembly, the World Summit on the Information Society, the Internet Governance Forum and the UN General Assembly, to name a few. Other issue areas are coming forward in these venues, including: promoting IPv6 deployment and advancing connectivity by promoting IXPs; advancing DNS SECURITY; generating a road map for future evolution of Internet governance; providing reliable tools for e-commerce, banking, private communications, etc., to move toward a more secure Internet; establishing work programs and guidelines for defining telecommunication development questions and priorities; and identifying properties for global Internet cooperation.

POLITICAL INTERESTS

Governments will always play a huge part in solving big problems. They set public policy and are uniquely able to provide the resources to make sure solutions reach everyone who needs them. They also fund basic research, which is a crucial component of the innovation that improves life for everyone.

– Goldstein (2010)

The fourth strategic area of competition is political. The Internet has become a political platform for messaging. Political actors now have the opportunity to perform on a global stage and compete to persuade multiple audiences at the same time, articulating policies and investments needed for strength in, and dominance of, the digital economy and that ultimately serve their own interests. They articulate the benefits quite clearly in terms of GDP growth, job creation, access to information and the ability to innovate. They also communicate the challenges in terms of threats to society, and the need to prepare for action and defend critical infrastructures, services, businesses and citizens from malicious cyber activities. With each speech given or initiative carried out, they position themselves for economic, political and military leverage, power and dominance.

As political actors communicate with their citizens — the constituency that holds the key to their power, legitimate or not — they highlight the rights of the individual to Internet access, better education, employment opportunity, economic well-being and privacy. When speaking to industry and government leaders, they highlight the need for partnership, emphasizing the link between delivery of citizen-essential services and state responsibility (in the manner by which the state dictates and by which a company can make a profit). But does their success in arguing for such deep partnership mean that a specific industry is working for national economic, political and military interests? Sub-rosa messages are also being conveyed, but the question of what market levers a state needs to impose to ensure collective market dominance and hence mutual economic growth remains.

Finally, some leaders are signalling thresholds and trying to establish norms of acceptable behaviour for other leaders.⁷ Their intent is to protect the value of their current and

future digital investments and to preserve the importance of the Internet for their political and economic interests. For example, President Xi Jinping has openly announced China's dual focus on developing technology and ensuring cyber security. These two aspects, he asserted, are "two wings of a bird" and require an overall plan to advance both simultaneously (Tiezzi 2014). Chancellor Angela Merkel has stated that "a 'cyber dialogue' is needed to set mutual privacy standards and legal frameworks...to catch up to rapidly advancing technology" (CBS 2014). Russian President Vladimir Putin has discussed similar governance issues, stating that "establishing international control over the Internet using the monitoring and supervisory capabilities of the [ITU]...[should be a]...priority on the international agenda" (Brito 2012). And US President Barack Obama has shared his viewpoint and concerns by stating that "America's economic prosperity, national security, and our individual liberties depend on our commitment to securing cyberspace and maintaining an open, interoperable, secure, and reliable Internet. Our critical infrastructure continues to be at risk from threats in cyberspace, and our economy is harmed by the theft of our intellectual property" (The White House 2014b).

Ultimately, the Internet remains both a global commons and part of each nation's sovereign infrastructure, and thus activities in cyberspace must continue to navigate two sets of demands: national interests and global interests. The forms of competition and tension discussed in this chapter are about different power struggles. They are also about those leaders who are using sophisticated strategies to forge complementary activities that ultimately serve their and their countries' interests. For those in the middle of this competition, it is important that they recognize that the gamesmanship and strategies are multifold. Perhaps this is why government intervention in this field tends to be more pronounced and pervasive — from controlling market access to subsidizing market entry and market share to imposing greater security requirements (and gaining access to intellectual property) to increasing censorship and surveillance practices for security and stability purposes. Political leaders are responsible for articulating a vision and establishing general principles and policies to achieve their goals and, accordingly, are constantly trying to advance their agendas using policy, law, market mechanisms, regulation, standards and other initiatives. The evidence is clear, you just have to look for it.

⁷ There are multiple venues where norms setting is taking place. The United Nations' Group of Governmental Experts, for example, facilitates dialogue among states to reduce risk and protect critical national and international infrastructure. It seeks consensus among nations on the applicability of the UN Charter, international law, and the principles of state sovereignty and responsible state behaviour to cyberspace. Additionally, the Organization for Security and Cooperation in Europe has been working on confidence-building measures to reduce the risks of conflict stemming from the use of information and communications technology. The report from that work was published in December 2013.

SOCIAL INTERESTS

Advances in the technology of telecommunications have proved an unambiguous threat to totalitarian regimes everywhere.

– Kotkin (2008)

The fifth and final strategic area of competition concerns the social aspects of the Internet and whether the Internet should be considered a citizen right or privilege. In less than two decades, the Internet has evolved from an opt-in service, where citizens and governments were able to choose whether or not to participate in the Internet society, to a compelled infrastructure that requires participation in order to reap its benefits and deliver essential services to citizens. This, in turn, is changing perceptions regarding citizens' rights and privileges. It is also shifting the power and perception of ownership.

In 2011, a group of nations formed the Freedom Online Coalition to advance Internet freedom — free expression, association, assembly and privacy online. During the 2014 NETmundial meeting, participants agreed that human rights should underpin Internet governance principles (NETmundial 2014b). Echoing the UN Human Rights Council's 2012 decision (United Nations 2012), they declared that the rights that people enjoy off-line must also be protected online in accordance with existing international human rights treaties and legal obligations.⁸ Some of these rights include freedom of expression, freedom of association, privacy, freedom of information and access to information. But if citizens really are to enjoy these rights, then what mechanisms do they have at their disposal to challenge their national leaders when their rights are violated? And who is going to enforce them? Unfortunately, the reality is that the very interconnectedness of people can be denied and that freedom of communication and political freedoms are clearly linked.

For example, many protests were organized in the *favelas* of Brazil leading up to the recent World Cup games. The citizens of historically underserved communities were angry over their living conditions and the government's pacification program, which, building on an earlier program, was "designed to seize back control of the areas from drug traffickers and make them safer for the tournament and the 2016 Olympics" (Bainbridge 2014). They were also angry about the amount of investment the government was making in the stadiums and facilities needed to support the influx of tourists during the World Cup, arguing that these resources would be better used

⁸ Including the International Covenants on Civil and Political Rights and Economic, Social and Cultural Rights, and the Convention on the Rights of Persons with Disabilities.

to improve the living conditions of its own citizens. In addition, they believed that their views were not represented by the quasi-state-controlled media and took matters into their own hands. Citizens became journalists — using their smartphones, digital cameras and apps such as Twitcast and Twitcam to circulate photos and videos so the world could see what was really happening in the streets of Brazil. Venezuela's government is also facing outraged citizens, and has blocked images on Twitter after violent protests emerged in Caracas seeking redress for "a catalogue of woes that include rampant inflation, food shortages and one of the world's highest murder rates" (Bajak 2014).

A related question is whether the "governed" have a right to own their data or to know what their "governors" (which can include both governments and private actors) are doing with their data. In May 2014, the European Court of Justice (ECJ) ruled in favour of a Spanish citizen's rights to privacy and sent a message to the data aggregators and content brokers that privacy is paramount. The ECJ's ruling upholds "European citizens' 'right to be forgotten,' that is, their right to have embarrassing and currently misleading information deleted from the Internet" (Farrell 2014). Many Europeans celebrated the ECJ ruling against Google, noting that the United States has not curbed the monopolistic behaviour of Google and its broad infringement on the privacy of citizens. For some, the ECJ's ruling was Europe's way of mitigating such behaviour. Governments are also believed to be infringing upon citizens' right to privacy. To address this concern, the US National Security Agency's Internet surveillance programs are being scrutinized, and President Obama recently pared down the scope of its collection activities (The White House 2014c). The United Kingdom and many other Western nations are also reviewing the scope of their intelligence services and some leaders are calling for new laws to govern surveillance programs (Ashford 2014).

On the other hand, other countries are supplementing their own surveillance practices by passing laws to require that data be stored within their territories, making it easier to intercept, search or protect. For example, "Russia's Parliament has approved a law similar to China's that would require Internet companies such as Google to locate servers handling Russian traffic inside the country and store user data locally for six months" (Khrennikov and Ustinova 2014).

Finally, when do the empowered go too far? Governments are increasingly requesting and can even compel private sector assistance in conducting voice or data surveillance. In some cases, there is no territorial limitation on that power. For example, Microsoft is fighting a US government search warrant that compels Microsoft to hand over customer data (emails) maintained in a data centre operated by one of its subsidiaries located in Ireland. The data in Microsoft Ireland's possession, custody or control

relates to a drug investigation (Nakashima 2014).⁹ This type of overt collection and government intervention is compromising the integrity of multinational companies that provision Internet services and store customer data. It also is contrary to and undermines existing international law. Many countries — including the United Kingdom, India, Belgium and the United Arab Emirates — are passing legislation to compel companies to hand over encryption keys to aid law enforcement investigations and support national security matters. Still others, China among them, are demanding that companies that want to deliver products to their (broadly defined) national security marketplace must turn over the source code for their products. More recently, perhaps in an effort to limit market penetration, a leading Chinese news agency branded Microsoft's Windows 8 operating system as a threat to the nation's information security (Williams 2014).

In the next five years, the number of global Internet users will double. That growth will primarily come from China, India and African nations. Those societies have very different histories, development trajectories, cultural backgrounds and experiences with government. Freedom of expression may not have the same cultural undertones (and support) as it has in the West. And experience in other areas shows that guaranteeing freedom of, and access to, information can be difficult, even if the necessary legislation is in place. How these new Internet users assert their voice, leverage their market positioning as consumers and influence power will show us whether they see the Internet as a citizen right or a privilege.

CONCLUSION

Two roads diverged in a wood, and I — I took the one less traveled by, and that has made all the difference.

– Frost (1979)

We are in the midst of an intense competition for money, power and control over all aspects of the Internet and the Internet economy. The competition for Internet dominance is being waged across economic, technical, regulatory, political and social battlefields. The web of relationships between each issue is noteworthy to say the least.

Underpinning this competition is the perception that the United States remains the Internet's superpower, a perception that many around the world would like to see change. The continuous release of information over the past year about the US government's role in Internet

surveillance and intervention has accelerated national desires and agendas to transfer Internet governance to venues such as the United Nations, the ITU and other international fora, which many perceive to be more legitimate, fair and transparent. Countries arguing for these significant changes are already establishing their own foothold on Internet matters, while also eroding the positions of the United States (and the West). This situation is also giving rise to private companies that feel violated by their own governments and are losing real market share around the world as a consequence.

Looking to a future where the demand curve and market growth of the Internet are likely to be driven from Asia, Africa and South America, the United States will not maintain its position of influence unless it develops and delivers a new message focused on economic competitiveness and business opportunity that respects the rights of individuals in their liberty, thoughts and possessions. Without a new cadre of leaders — both in the government and in the private sector — it will be very difficult for the United States to engage around the globe without being perceived as colonialist or paternalistic. And the chorus calling for multilateral organizations to seize control over the technical and regulatory underpinnings of the Internet will only continue to grow in volume and power.

Counteracting these calls for change requires a new message that can unify nations in a common vision of how the Internet and its underlying technologies can foster trust, fuel global economic growth for all and empower citizens. A thorough action plan that brings together a broad set of countries and participants to work toward this vision — jointly and across borders, and in partnership with government and non-state actors — is the way forward.

Who will stand up and be the guarantor of the Internet's future? America's strategic interests are at stake and, as David versus Goliath, the world is now rooting for David to win.

⁹ On July 31, 2014, Judge Loretta A. Preska of the United States District Court for the Southern District of New York upheld the position of the US Government; however, the court granted a stay of its decision pending appeal to enable Microsoft to appeal. See Joseph Falcone, "US Federal Court Orders Microsoft to Produce E-Mail Content Stored Outside the United States," Herbert Smith Freehills New York LLP, 2014.

WORKS CITED

- Ashford, Warwick. 2014. "Most NSA Spy Data Relates to Innocent Internet Users, Report Shows." *ComputerWeekly.com*, July 7. www.computerweekly.com/news/2240223999/Most-NSA-spy-data-relates-to-innocent-internet-users-report-shows.
- Bainbridge, Luke. 2014. "How Social Media Gives New Voice to Brazil's Protest." *The Guardian*, April 26. www.theguardian.com/world/2014/apr/27/social-media-gives-new-voice-to-brazil-protests.
- Bajak, Frank. 2014. "Venezuela Cuts off Internet, Blocks Communications for Protestors." *The Huffington Post*, February 21. www.huffingtonpost.com/2014/02/21/venezuela-internet_n_4832505.html.
- Bilbao-Osorio, Beñat, Soumitra Dutta and Bruno Lanvin, eds. 2014. *The Global Information Technology Report 2014: Rewards and Risks of Big Data*. World Economic Forum. www3.weforum.org/docs/WEF_GlobalInformationTechnology_Report_2014.pdf.
- Bradner, S. and A. Mankin. 1995. "The Recommendation for the IP Next Generation Protocol." RFC 1752. January.
- Brito, Jerry. 2012. "The Case Against Letting the U.N. Govern the Internet." *Time*, February 13. <http://techland.time.com/2012/02/13/the-case-against-letting-the-united-nations-govern-the-internet/>.
- CBS. 2014. "Merkel: 'Difficulties Yet To Overcome' In US Spy Scandal." <http://washington.cbslocal.com/2014/05/02/merkel-difficulties-yet-to-overcome-in-us-spy-scandal/>.
- CIGI. 2015. "Global Commission on Internet Governance." www.cigionline.org/activity/global-commission-internet-governance.
- Deutsche Welle. 2014. "Weighing a Schengen Zone for Europe's Internet Data." *Deutsche Welle*, February 20. www.dw.de/weighing-a-schengen-zone-for-europes-internet-data/a-17443482.
- Euractiv. 2014. "EU Internet Governance: Franco-German Alliance." *EurActiv.Com*, July 11. www.euractiv.com/sections/innovation-enterprise/eu-internet-governance-franco-german-alliance-303421.
- European Commission. 2013. "Proposal for a Directive of the European Parliament and of the Council: Concerning Measures to Ensure a High Common Level of Network and Information Security across the Union (COM(2013) 48 final)." Brussels: European Commission.
- Evans, Dave. 2011. *The Internet of Things. How the Next Evolution of the Internet Is Changing Everything*. San Jose, CA: Cisco Internet Business Solutions Group. www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf.
- Evans, Peter C. and Marco Annunziata. 2012. *Industrial Internet: Pushing the Boundaries of Minds and Machines*. General Electric. November 26.
- Farrell, Henry. 2014. "Five Questions About the European Court of Justice's Google Decision." *The Washington Post*, May 14. www.washingtonpost.com/blogs/monkey-cage/wp/2014/05/14/five-key-questions-about-the-european-court-of-justices-google-decision/.
- Fitchard, Kevin. 2013. "Politics Could Pop Google's Project Loon." *Business Week*, June 24. www.businessweek.com/articles/2013-06-24/politics-could-pop-googles-project-loon.
- Frost, Robert. 1979. "The Road Not Taken." In *The Poetry of Robert Frost: The Collected Poems, Complete and Unabridged*.
- Goldstein, Dana. 2010. "5 Questions for Bill Gates: The Full Interview." *Daily Beast*, January 24. www.thedailybeast.com/articles/2010/01/25/5-questions-for-bill-gates-the-full-interview.html.
- Hathaway, Melissa E. 2010. "Toward a Closer Digital Alliance." *SAIS Review* 30 (2).
- Hathaway, Melissa E. and John E. Savage. 2012. "Stewardship of Cyberspace: Duties for Internet Service Providers." *Cyber Dialogue* 2012. March. http://belfercenter.ksg.harvard.edu/files/cyberdialogue2012_hathaway-savage.pdf.
- Hathaway, Melissa E. and Alexander Klimburg. 2012. "Preliminary Considerations: On National Cyber Security." Chapter 1 in *National Cyber Security Framework Manual*, Tallinn, Estonia: NATO Cooperative Cyber Defence Centre of Excellence, December 2012.
- IETF. 2015. "The Internet Standards Process." Revision 3. www.ietf.org/about/standards-process.html.
- Internet World Stats. 2014. "Internet Penetration by Region." www.internetworldstats.com/stats.htm.
- ITU. n.d. "The Constitution of the International Telecommunication Union, Preamble."
- . 1989. "Final Acts of the World Administrative Telegraph and Telephone Conference Melbourne, 1988: International Telecommunications Regulations." Geneva, Switzerland. www.itu.int/dms_pub/itu-s/oth/02/01/S02010000214002PDFE.pdf.

- Jackson, Joab. 2013. "NIST Denies NSA Tampering with Encryption Standards." *PC World*, September 10. www.pcworld.com/article/2048510/nist-denies-nsa-tampering-with-encryption-standards.html.
- Keynes, John Maynard. 1935. "Chapter 21: The Theory of Prices." In *The General Theory of Employment, Interest, and Money*. www.marxists.org/reference/subject/economics/keynes/general-theory/ch21.htm.
- Khrennikov, Ilya and Anastasia Ustinova. 2014. "Putin's Next Invasion, the Russian Web." *Business Week*, May 1. www.businessweek.com/articles/2014-05-01/russia-moves-toward-china-style-internet-censorship.
- Kotkin, Stephen. 2008. "How Murdoch Got Lost in China." *The New York Times*, May 4. www.nytimes.com/2008/05/04/business/media/04shelf.html?_r=0.
- Levine, Yasha. 2014. "Putin Ramps up Internet censorship, citing Google and Snowden to ensure public Support." *Pando Monthly*, March 20. <http://pando.com/2014/03/20/putin-ramps-up-internet-censorship-citing-google-and-snowden-to-ensure-public-support/>.
- Nakashima, Ellen. 2014. "Microsoft Fights U.S. Search Warrant for Customer e-mails held in Overseas Server." *The Washington Post*, June 10. www.washingtonpost.com/world/national-security/microsoft-fights-us-search-warrant-for-customer-e-mails-held-in-overseas-server/2014/06/10/6b8416aef0a7-11e3-914c-1fbd0614e2d4_story.html.
- NETmundial. 2014a. "Global Stakeholder Meeting on the Future of Internet Governance." <http://netmundial.org>.
- . 2014b. "NetMundial Multi-stakeholder Statement." April 24. <http://netmundial.org/netmundial-multistakeholder-statement/>.
- Rash, Wayne. 2012. "WCIT Treaty Talks End in Dubai With Walkout of U.S. and Allies." *E-Week*, December 15. www.e-week.com/cloud/wcit-treaty-talks-end-in-dubai-with-walkout-of-us-allies/.
- Rediff Business. 2013. "20 Best Broadband Providers in the World." November 14. www.rediff.com/business/slide-show/slide-show-1-special-20-best-internet-broadband-providers-in-the-world/20131114.htm#2.
- SixXS. 2015. "IPv6 Transition Mechanism/Tunneling Comparison." www.sixxs.net/faq/connectivity/?faq=comparison.
- Smith, Adam. 1909. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Edited by C. J. Bullock. In *The Harvard Classics*, edited by Charles W. Eliot. New York, NY: P.F. Collier & Son.
- Sorav, Jain. 2012. "40 Most Popular Social Networking Sites Around the World." *Social Media Today*, October 6. www.socialmediatoday.com/content/40-most-popular-social-networking-sites-world.
- Sterling, Joe. 2012. "FBI Says it's Using Facebook, Twitter to Find 'Wanted Terrorist.'" *CNN*, October 3. www.cnn.com/2012/10/03/justice/massachusetts-fbi-terrorist/.
- Sun Tzu. 1963. *The Art of War*. New York: Oxford University Press.
- TechWatch. 2014. "Top 10 IT Data Aggregators." www.jazdtech.com/techdirect/leaf/Data-Management/Database-Tools/Data-Aggregation.htm.
- The Federal Trade Commission. 2014. "Data Brokers: A Call for Transparency and Accountability." May. www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf.
- The White House. 2013. "Improving Critical Infrastructure Cybersecurity." Executive Order 13636. February 12. www.gpo.gov/fdsys/pkg/FR-2013-02-19/pdf/2013-03915.pdf.
- . 2014a. "Big Data: Seizing Opportunities and Preserving Values." May 1. www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_5.1.14_final_print.pdf.
- . 2014b. "Statement by the President on the Cybersecurity Framework." February 12. www.whitehouse.gov/the-press-office/2014/02/12/statement-president-cybersecurity-framework.
- . 2014c. "Signals Intelligence Activities." Presidential Policy Directive 28. January 17. www.whitehouse.gov/sites/default/files/docs/2014sigint_mem_ppd_rel.pdf.
- The Tor Project. 2015a. "Censorship." <https://blog.torproject.org/category/tags/censorship>.
- . 2015b. "About." www.torproject.org/about/overview.html.en.
- Tiezzi, Shannon. 2014. "Xi Jinping Leads China's New Internet Security Group." *The Diplomat*, February 28. <http://thediplomat.com/2014/02/xi-jinping-leads-chinas-new-internet-security-group/>.

Troianovski, Anton and Danny Yardin. 2014. "German Government Ends Verizon Contract." *The Wall Street Journal*, June 26. <http://online.wsj.com/articles/german-government-ends-verizon-contract-1403802226>.

United Nations. 2012. United Nations "The Promotion, Protection and Enjoyment of Human Rights on the Internet." General Assembly, Human Rights Council Twentieth Session, 20/L13. June 29. Office of the High Commissioner for Human Rights. www.loc.gov/lawweb/servlet/lloc_news?disp3_l205403231_text.

US Department of Commerce. 2014. "NTIA Announces Intent to Transition Key Internet Domain Name Functions." March 14. www.ntia.doc.gov/press-release/2014/ntia-announces-intent-transition-key-internet-domain-name-functions.

Williams, Wayne. 2014. "China Brands Windows 8 a Threat to Its National Security." *Beta News*, June 5. <http://betanews.com/2014/06/05/china-brands-windows-8-a-threat-to-its-national-security/>.

Zmijewski, Earl. 2014. "Turkish Internet Censorship Takes a New Turn." *ReSys*, March 30. www.renysys.com/2014/03/turkish-internet-censorship/.

ABOUT THE AUTHOR

Melissa E. Hathaway is a CIGI distinguished fellow and is contributing to the Global Security & Politics Program's research on Internet governance.

Melissa is president of Hathaway Global Strategies LLC, where she brings a multidisciplinary and multi-institutional perspective to strategic consulting and strategy formulation for public and private sector clients. Additionally, Melissa is an appointed member of the Global Commission on Internet Governance and a senior adviser at Harvard Kennedy School's Belfer Center.

Melissa served in two presidential administrations, where she spearheaded the Cyberspace Policy Review for President Barack Obama and led the Comprehensive National Cybersecurity Initiative for President George W. Bush. At the conclusion of her government service, she received the National Intelligence Reform Medal and the National Intelligence Meritorious Unit Citation Medal in recognition of her achievements.

Melissa has a B.A. from The American University in Washington, DC, and is a graduate of the US Armed Forces Staff College, with a special certificate in information operations.

**CHAPTER SIX:
INCREASING INTERNET CONNECTIVITY WHILE COMBATting
CYBERCRIME: GHANA AS A CASE STUDY**

Caroline Baylon and Albert Antwi-Boasiako

Copyright © 2016 by Caroline Baylon and Albert Antwi-Boasiako

ACRONYMS

CET	Common External Tariff
CID	Criminal Investigations Department
DDoS	distributed denial of service
ECOWAS	Economic Community of West African States
ISPs	Internet service providers
IT	information technology
SAT-3/WASC	South Atlantic 3/West Africa Submarine Cable
USB	universal serial bus
VAT	Value Added Tax

INTRODUCTION

Despite the steady development of Internet infrastructure in recent years in Ghana, the country still faces considerable obstacles to achieving widespread connectivity. Ghana's economy has been growing consistently, and measures to further foster Internet penetration are central to Ghana's continued economic development. But along with the Internet's tremendous benefits for business and commerce comes a challenge: greater Internet connectivity in the country is correlated with an increase in cybercrime. Ghana is already one of the top 10 sources of cybercrime in the world, and faster and more reliable Internet will provide cybercriminals with greater opportunities to engage in illicit activities online — and with a larger number of potential victims (Sikiti da Silva 2014). This chapter looks at how to best help foster Internet growth in Ghana while simultaneously working to contain cybercrime levels.

The first section of this chapter focuses on Internet infrastructure development, examining the current state (and evolution) of Internet connectivity in Ghana, the root causes of the challenges to Internet infrastructure development and potential solutions to these challenges. The second section elaborates on the link between growth in Internet connectivity and increases in cybercrime. The third section centres on combatting cybercrime, looking at the current state (and evolution) of cybercrime in Ghana, the root causes of cybercrime and potential solutions to these challenges. The fourth section considers whether policy makers can develop an overarching strategy to tackle these interlinked problems. The chapter draws on the existing academic literature, local news sources in Ghana and data gathered by reports made to the e-Crime Bureau.

INTERNET INFRASTRUCTURE DEVELOPMENT

The Current State and Evolution of Internet Infrastructure Development in Ghana: A Snapshot

The First Wave of Connectivity: Internet Cafés

Ghana was one of the earliest countries in Africa to gain Internet connectivity in 1994 (Foster et al. 2004). The first wave of connectivity in the country came through fixed-line access. Until 2010, Ghana had only one submarine fibre optic cable, the SAT-3/WASC (South Atlantic 3/West Africa Submarine Cable), which it shared with other West African countries. With such a limited fixed-line network, the primary method of Internet access was shared connectivity, mostly in Internet cafés. This also included other public access places such as the workplace, schools and universities. Fixed-line connectivity in homes remains rare: only three percent of households in Ghana had a working Internet connection in 2012 (Calandro, Stork and Gillwald 2012).

Poor Quality, an Urban-Rural Digital Divide and Last-mile Connectivity Issues

Moreover, connectivity in the country is often of poor quality. The network experiences frequent outages as well as slow Internet speeds. According to a 2012 survey, more than 40 percent of respondents reported that “slow Internet” limited their use (Frempong 2012).

Ghana also has a significant urban-rural digital divide, with the majority of the country's Internet connectivity (and especially faster fibre optic connections) concentrated in the capital city of Accra and other large cities. Rural areas often lack sufficient economic incentives for investment: wages are lower and, since they tend to be agricultural communities, there is less market demand for connectivity. The cost of providing Internet to rural areas is significantly higher as well. They typically lack last-mile infrastructure and thus may need satellite or other connections, which are expensive. Rural areas are less likely to have electricity, compounding the challenge.

The Second Wave of Connectivity: Mobile Internet

Given the challenges involved in providing Internet in the region, Ghana has increasingly turned to mobile broadband, which has formed the second wave of connectivity in the country. Smartphones and access dongles (i.e., universal serial bus [USB] modems with a SIM card inside that can be plugged into a computer) are now the main means of Internet access in Ghana. The country's mobile broadband penetration rate recently reached 62 percent (National Communications Authority 2015). Internet cafés, once the

primarily place for Ghanaians to access the Internet, have been declining in popularity (Acquaye 2013).

The rise of mobile broadband is in large part due to its increased affordability. While fixed-line Internet in Ghana requires a subscription, mobile Internet in the country is primarily based on prepaid services. (More than 97 percent of mobile phone owners in Ghana are on a prepaid plan.) Since many Ghanaians hesitate to sign up for subscription-based services out of concern that they will not be able to pay in subsequent months, the prepaid nature of mobile broadband is much more attractive for them (Calandro, Stork and Gillwald 2012). And for those accessing the Internet via smartphone, mobile phones are much less expensive than computers as well.

Another advantage of mobile Internet is that it does not require users to have electricity at home. Of course, mobile phones need charging, but they can be charged at regular intervals at a venue where there is electricity and then taken elsewhere. Mobile broadband is also helping to bridge the digital divide by bypassing the last-mile connectivity problem — although many villages are still grappling with the challenges stemming from lack of electricity.

However, a number of obstacles still remain. The experience of viewing the Internet via smartphone is not akin to accessing the Internet from a computer; the devices' small screen size and limited computing power mean that users cannot access information as readily. And with access dongles, the coverage is often spotty; nor do they provide enough capacity. It is thus important for Ghana to continue to develop its fixed-line connectivity too.

The Coming Third Wave? Recent Developments in Fixed-line Connectivity

In the past five years, Ghana has acquired four additional submarine fibre optic cables: Main One Cable in 2010, GLO-1 (Globacom-1) and WACS (West Africa Cable System) in 2011, and ACE (African Coast to Europe) in 2013. This has considerably increased the bandwidth available, from 320 gigabytes to 12 terabytes. It has also resulted in “a dramatic fall in the wholesale cost of capacity. Today, the cost of an E1 connection in Ghana is around \$1,200, down from as much as \$12,000 in 2006” (Boakye 2014). That is, the cost for Internet service providers (ISPs) to purchase bandwidth is now roughly one-tenth of what it was less than 10 years ago. These cost savings are starting to be passed on to consumers, although this has not yet been fully accomplished.

As part of its Project Link, Google announced in October 2015 that it plans to lay 2,000 km of fibre in Accra and the major cities of Tema and Kumasi (Abdul-Jalil 2015). Meanwhile, the government is taking steps to enhance Internet connectivity in rural areas: it recently finished building an 800-km fibre optic cable backbone traversing

the eastern corridor of the country and is starting a similar project for the western corridor (IT News Africa 2015; Acquaye 2015).

Challenges for Internet Infrastructure Development: Identifying the Root Causes

Cost Factors, including High Poverty Levels

Cost factors have long precluded the greater development of the Internet in Ghana, especially fixed-line connectivity. Given that 54 percent of the population lives on less than GH¢7 (US\$2) per day, computers remain unaffordable for the majority, and only 9 percent of households in Ghana had a computer in 2012 (Dela Klutse 2015; Calandro, Stork and Gillwald 2012). For many, the cost of Internet access alone is prohibitive; over 55 percent of those interviewed in a 2012 survey said that their main reason for not accessing the Internet was because it was “expensive to use” (Frempong 2012). Concerns about the size of the market have thus made some telecommunications companies reticent to invest the tens or hundreds of millions of dollars needed for laying down cables, including in rural areas.

Electricity Shortages

The insufficiency and unreliability of the electricity supply, even in Accra, is a major challenge for Internet development too. Only 73 percent of households in Ghana had electricity in 2012, and those that do experience frequent power cuts (Calandro, Stork and Gillwald 2012). In December 2014, for example, the capital experienced weekly blackouts that lasted for up to 12 hours at a time.

This is because much of Ghana's energy comes from hydroelectric power. The country's rainfall is unpredictable, so at times the lake that supplies the country's main hydroelectric power plant does not contain enough water. These electricity shortages are a major obstacle to the use of computers needed for fixed-line Internet access. It also raises the cost of providing Internet, since ISPs must have electrical or diesel generators or other backup methods to provide power when there are cuts.

Accidental Cable Cuts during Road Construction and Repairs or Illegal Mining

Ghana also has a major problem with unintentional cuts to both fibre optic and copper cables, resulting in poor quality and outages that can last days or even weeks. In 2014, 1,370 fibre optic cable cuts were reported in a six-month period, or more than 200 cuts per month. This represents a 400 percent increase since 2011, in which there were only 480 cuts total for the year, an average of 40 cuts per month (Naphtal 2015). The cable cuts result in considerable expense for telecommunications companies, in both the cost of replacing the cables and the labour required to repair them. Repairing one cut

costs an estimated GH¢17,000 (US\$4,500). Moreover, the labour spent repairing the cables could have been used to improve Internet service or lay additional cables instead (Kunateh 2015).

Repairs to fibre optic cables involve slow and delicate work. Just identifying the location of a cable cut can take five to six hours. The fibre optic cables themselves are very sensitive: each cable contains more than 46 glass fibre strands and each strand has to be cut to the right shape in order to be spliced back together. Moreover, if the protective coating has been damaged, the cables must be carefully cleaned of dust particles before being reconnected; otherwise, it will create interference on the line. Another complication is that if a fibre optic cable experiences too many cuts, this eventually causes “attenuation.” That is, the signal will meet resistance when it passes down the fibre, degrading the communication quality. Repeated cuts may require telecommunications companies to replace a whole section of cable, which is costly (Adam 2014).

The majority — an estimated 75 percent — of cuts are caused by workers carrying out road construction and repairs (BiztechAfrica 2013). Ironically, many of these workers are contractors of the Ministry of Roads and Highways, which has been actively improving and expanding the road network. Although recent efforts have been made to ensure that the locations of the cables are signposted and that blueprints are filed with relevant agencies such as the Ghana Highway Authority, some managers do not ask for the blueprints. Moreover, the managers generally do not explain the importance of the cables to the workers or make sure that they know how to recognize the signposts. Instead, they give them key performance indicators, so the workers rush to complete their tasks without considering the cables (Graphic Online 2012). The expansions of some roads from one lane to two may also disrupt the existing cable infrastructure. In other instances, the cables may not have been buried deep enough or the blueprints may be unclear. Despite the numerous reports that telecommunications companies have made to the Ministry of Roads and Highways, this problem persists (Mustapha 2014).

Other culprits include contractors installing road signs or working for utility companies to lay water pipes and electrical lines. The contractors sometimes cut through roads in the process, damaging the roads, which then require further repair, further perpetuating the cycle of cable damage (Ghana Business News 2015). About 10 percent of cable cuts are caused by illegal miners, or “galamsayers,” who are digging for minerals in the ground (Acquaye 2014). They, too, often do not understand the importance of the cables.

Theft of Cables and Other Elements of the Internet Infrastructure

Cable theft is a major challenge. The relatively high price of copper means that thieves are increasingly digging up and stealing copper cables, then selling them to scrap dealers who resell them abroad. In January 2013, Vodafone experienced the theft of 1 km of copper cable in the Madina area, which was estimated to cost GH¢200,000 (US\$53,000) in damages and to take three weeks to repair. And in the month of May 2013 alone, Vodafone experienced about 30 cable thefts, with outages affecting Osu Castle (the seat of government at the time) and Parliament House in Accra, the capital (JOY Online 2013).

The theft of copper cables also contributes to cuts to (and sometimes the theft of) fibre optic cables. While costing millions to make, fibre optic cables fetch little on the black market. They are primarily made of glass, and thus have limited scrap resale value. When purchased, the glass is primarily reused to make jewellery. However, thieves looking for copper cables sometimes find the fibre optic cables and think that they might contain copper, so they cut them open to check. On discovering that they do not, the thieves sometimes steal the fibre optic cables anyway, thinking that even the small amount of money they will fetch is better than none.

ISPs also have to guard against the theft of other elements of the infrastructure required to provide Internet. For example, there have been numerous incidences of theft of the diesel from the diesel generators used to provide backup power. Similarly, the telecommunications provider Tigo has seen an increase in thefts of batteries, which are also used as a backup power source, and has had to increase security in response (Ghana News Agency 2014). This entails additional costs, both in terms of replacing materials and providing security, and thus also raises the cost of providing Internet.

Corruption

The high incidence of corruption within the country is also a factor in cable theft and impedes Internet development. Corruption is pervasive even within major companies. In one instance, four employees of Vodafone — including a team manager and a senior customer engineer — were jailed for stealing cables from the company. The company had sent the employees to recuperate some 200 m of redundant Vodafone underground cables for use in repair works, but instead the employees took almost seven times that amount and sold the extra to a scrap dealer. When those within a company, who — in the words of the trial judge — “ought to protect the company and not engage in such acts that would make the company run at a loss” are corrupt, this compounds the challenge of further developing Internet infrastructure (GhanaWeb 2015a).

Potential Solutions

Tax Incentives

One method of increasing Internet penetration is to lower or eliminate taxes on equipment needed for access, which would make it more affordable for individuals. Doing so for smartphones may prove effective. Ghana has employed such tactics with mobile phones (not just smartphones) in the past: it removed import duties on all mobile phones in 2008. Although the government later reintroduced tariffs in 2013 due to its need to raise additional funds at the time, the increase in uptake of mobile phones during the tax-free period demonstrates the success of the policy (Citifmonline 2016b).

With this in mind, Ghana has recently reduced taxes on smartphones specifically. It had initially intended to repeal all of the customs charges that had been placed on smartphones: in November 2014, the country's finance minister announced that he planned to remove import duties on smartphone handsets, stating that this would bolster smartphone penetration rates and thus also help close the digital divide (Ogundeji 2014). He also expressed his view that, despite the loss in revenue from import duties, the measure would result in an increase in overall tax revenue; a reduction in the smuggling of handsets into the country combined with an increase in the number of smartphones sold would increase revenue from other taxes, such as the Communication Service Tax, Value Added Tax (VAT), and corporate taxes.

However, the tariff removal never went into effect since it subsequently became apparent that, as a member of the Economic Community of West African States (ECOWAS), Ghana was expected to implement the Common External Tariff (CET) when it came into force in early 2016 (Citifmonline 2016a). Instead, the government has therefore recently reduced the customs tariffs on handsets from 20 percent to 10 percent, so that they are in line with the CET. It also removed the VAT on imported handsets. While not as effective a measure as removing taxes on smartphones entirely, the tax reduction should nonetheless contribute to a significant increase in smartphone penetration and hence in mobile Internet connectivity.

The government should consider repealing or lowering import duties on computers too. There is currently an import tax exemption on computers used for educational purposes. However, broadening this to include all computers would help promote fixed-line access.

Another effective tactic would be to reduce or remove taxes on Internet infrastructure equipment, making investment more affordable for companies. Telecommunications companies are petitioning the government to remove taxes on modems (also called terminal equipment) used to access the Internet, including access dongles. This will

reduce the cost for these companies to provide Internet. They point out that the tax reduction on smartphones favours the provision of Internet via mobile. Since each type of Internet access has important advantages and both are needed, reducing taxes on modems and computers as well as on smartphones would stimulate both fixed-line and mobile connectivity.

Renewable Energy

To increase the reliability of the energy supply, greater development of solar power (and other renewables such as wind and biomass power) could serve as an ideal complement to hydro power. The current challenge in Ghana is a lack of expertise on how to implement renewable energy. However, there have been some promising steps in recent years. In November 2015, the Energy Commission organized a first conference on renewable energy that brought together key stakeholders including the private sector, financiers, government and consumers. It held a second conference in August of this year together with the United Nations Development Programme and the Ministry of Power.

Engaging All Stakeholders in Preventing Accidental Cable Cuts and in Fighting Cable Theft

Solutions to accidental cable cuts will require the cooperation of all stakeholders: government ministries, contractors, workers and telecommunications companies. Positive measures have been taken, but are far from being adequately implemented. In March 2013, the Ghana Chamber of Telecommunications and the Association of Road Contractors agreed to jointly engage in a sensitization program to ensure that workers could recognize and protect cables and other telecommunication infrastructure when they saw it in the field, that telecommunications companies provide the latest blueprints to road agencies and that road agencies pass them on to their contractors. A May 2013 letter from the National Security Council Secretariat to the minister of roads and highways called for contractors to be given cable blueprints and for them to sign an agreement to protect cables from damage, with copies of all documents sent to the National Security Council Secretariat. The minister of roads and highways consequently issued a directive in December 2014 requiring contractors to engage with telecommunications companies before doing road work and for them to share blueprints and technical plans. The Ministry of Roads and Highways plans to establish a Standing Technical Committee (made up of the Telecoms Chamber, the National Communications Authority and the ministry itself) to act as an advisory body to help identify cable markings before work on a road begins.

Solutions to cable theft will also require the government, police, members of the community and telecommunications companies to cooperate. In October 2012, community

volunteers undertook to patrol the town of Ashaiman after a spate of cable thefts; they caught two individuals trying to steal cables and turned them in to police (GhanaWeb 2012a). In June 2013, the Ministry of Communications announced that, working together with Vodafone and other telecommunications providers, it was implementing a year-long awareness program to educate the population about the consequences of cable thefts. Vodafone also launched a community vigilance campaign, working closely with the police. The campaign underlines the key role of the community in stopping cable thefts, and Vodafone provides a dedicated phone number for citizens to report any suspicious activity near cables (GhanaWeb 2012b).

Technical Solutions

ISPs are also looking at technical solutions. For instance, Tigo began a project mounting fibre optic cables on concrete poles in the Western and Ashanti regions in November 2014 (CRU Wire and Cable News 2014). Similarly, telecommunications provider MTN has built in redundancy to existing routes to mitigate the effects of cable cuts (JOY Online 2014). Some telecommunications companies have suggested to the World Bank and other funders of road projects that the projects should include utility ducts so that companies can lay cables in a manner that reduces the risk of cuts. Other technical solutions might involve alarm systems when cable lines are disturbed.

Regulatory Measures

Regulatory solutions are also being considered. In November 2015, one of MainOne's senior executives called on the government to pass laws to protect undersea cables, given their key role in delivering Internet to the country. Such a law would complement an existing industry initiative on the topic, which includes an annual Cable and Pipeline Protection Awareness Workshop founded by MainOne to raise the issue's profile with relevant stakeholders (GhanaWeb 2015b).

Better enforcement of existing laws is also key. In April 2016, at the urging of Vodafone, the Ghana Police Service and the Judicial Services department set up special "cable courts," designed specifically for prosecuting cases of cable destruction and theft. The cable courts are currently established in Accra and Kumasi, but the intention is to establish additional such courts in other parts of the country as well (Ampomah 2016; Abbey 2016). In addition, in August 2016 the Ministry of Trade and Industries granted permission to Vodafone to inspect all scrap exports out of the country to ensure that they do not contain copper cables (GhanaWeb 2016; Ghana News Agency 2016).

THE LINK BETWEEN INTERNET INFRASTRUCTURE DEVELOPMENT AND CYBERCRIME

Greater Resources and Lower Barriers to Entry: Email Scams and Crimeware

Yet as Ghana's Internet infrastructure continues to expand, more widespread, cheaper and faster Internet connectivity is giving cybercriminals greater resources to engage in illegal activity, including providing them with access to a larger number of potential victims. The ability to send a large number of emails to a global pool rapidly and without any postage costs enables Ghanaian cybercriminals to more effectively engage in email scams, targeting both victims in Ghana and comparatively wealthy foreigners abroad.

Moreover, there are relatively low barriers to entry to commit cybercrime — and they are getting lower still (Kavanagh 2013). Although more than 80 indigenous languages are spoken in Ghana, English is the official language and the lingua franca. This means that a number of cybercriminals in the country have English language skills that they can employ in email scams that target the large part of the world population that speaks English — notably individuals living in wealthy countries such as the United States and United Kingdom. Further, the rise in crimeware, or malware designed to automate and facilitate cybercrime, means that engaging in cybercrime involves less and less technical skill. For instance, cybercriminals can use exploit kits with pre-written exploit code that do not require expertise to use.

High Vulnerability to Attack: User Inexperience and Underprotected Machines

As increasing numbers of computers in Ghana connect to the Internet, they form a mass of vulnerable machines that are particularly attractive targets for cybercriminals. Many computers in Ghana are not patched regularly and do not have antivirus software installed. This is partly because of a lack of user education: users are typically unaware of the importance of downloading update patches or running antivirus, so their machines often lack basic security protections.

Even when users are aware of the need for updates and antivirus, slow connection speeds and limited bandwidth are an obstacle to installing them. Another factor is that many users in Ghana use pirated software, which means that their software does not receive automatic security patches in response to newly discovered vulnerabilities. (In contrast, genuine software automatically receives updates to install.) Low income levels in Ghana are a key part of the challenge, as many users cannot afford genuine software. In addition, users often do not have access to antivirus software in their native language, making it

harder for them to use. Even developing versions for a few of the most commonly spoken languages — for example, Akan, Ewe and Ga — may not be economically viable.

Botnets

Cybercriminals (both within Ghana and abroad) can infect these unprotected or underprotected machines and herd them together into botnets, or “robot networks” of tens or hundreds of thousands of compromised computers that they can remotely control. By harnessing the combined power of the computers in a botnet, cybercriminals can launch distributed denial of service (DDoS) attacks to take down websites and other targets by directing a large volume of traffic against them in order to overwhelm them. Cybercriminals can also use such botnets to send malware to infect other computers in order to steal passwords, log-in credentials, bank details, credit card information and other data from victims both within Ghana and around the world.

Mobile Phone Vulnerabilities

Ghanaians’ heavy reliance on cell phones for Internet access also renders them especially vulnerable given that mobile devices have fewer cyber security protections than computers. They typically do not have defensive measures such as firewalls, antivirus software and encryption. Mobile phones’ operating systems are also not updated as frequently as those on personal computers.

COMBATting CYBERCRIME

The Current State and Evolution of Cybercrime in Ghana: A Snapshot

Initial Focus on Scamming: Nigerian Letter Hoaxes, Business Fraud, Credit Card Fraud and Romance Fraud

The earliest form of cybercrime in Ghana, known as “sakawa,” focused on scamming attacks. Of the 217 incidents reported to the e-Crime Bureau for investigations assistance in 2014, scamming attacks accounted for nearly 65 incidents — the largest category. One of the oldest tactics is advance fee fraud, in which a fraudster promises a victim a large sum of money (which never materializes) in exchange for a smaller upfront payment. Originally perpetuated via postal mail, such scams began to be conducted through email once the country gained Internet connectivity in 1994. The most famous type of advance fee fraud is “Nigerian letter hoaxes,” thus named because of the country in which they originated. (They are also sometimes referred to as “419” scams after the section of the Nigerian penal code that they violate.) For instance, a typical scam might involve a fake inheritance scenario. A fraudster purporting to be a lawyer contacts a victim to tell them that they have inherited millions of dollars from a distant relative, but they need to pay taxes or other fees in

order to obtain the money; of course, no such inheritance actually exists.

Another type involves fake business transactions associated with the sale of gold. A scammer claiming to be from a gold mining company may contact a victim to offer them gold for sale at an advantageous price. The buyer must send money beforehand for shipping or other costs; of course, he will never receive any gold. Similar tactics exist involving the sale of oil at below market prices.

Credit card fraud is another early feature of electronic fraud in Ghana. Incidents first occurred in 1999, in which the staff in international hotel chains stole the credit card numbers of Western tourists and sold them to scammers. The scammers used the card numbers to make online purchases and have the goods shipped to Ghana (Warner 2011).

Identity fraud, in which a fraudster obtains and uses a victim’s personal data to make use of their identity for economic gain, is also common. Romance fraud is especially frequent. Typically, scammers use stolen photos to pose as Westerners on Internet dating sites such as Match.com or eHarmony. They develop romantic relationships with victims living abroad, then ask for money for an emergency, plane tickets to visit the victim or other fabricated reasons. Individuals can lose large sums of money: one recent UK victim of a romance fraud lost £250,000 (US\$330,000) (National Crime Agency 2014).

The Modus Operandi: Roots in Traditional Juju Beliefs and Crime Bosses

Those engaging in sakawa often believe that the use of juju — a term that encompasses a number of traditional West African religions involving the use of black magic or witchcraft — is essential for their scams to be successful¹ (Abubakar 2012). The premise of juju beliefs is that individuals can make payments to “mallams,” or priests, to bargain with the spirits on their behalf in order to acquire wealth or power (Morton 2011). When engaging in cybercrime, a number of scammers therefore consult mallams, who — in exchange for money, of course — give the scammers a series of rules and rituals to follow, such as wearing a magic ring or sleeping in a coffin alongside a corpse² (Warner 2011). If they carry out the mallam’s instructions, the scammers believe, they will be protected from being caught by the police and will also acquire the power to control the minds of their victims, who will send

1 Juju priests have been involved from the beginning, when these scams were perpetuated by postal mail. For instance, juju priests would “bless” letters sent to Westerners, the letter writers believing that the recipients would then be magically compelled to send larger amounts of money (Warner 2011).

2 In some instances juju may involve elements of child sacrifice and cannibalism.

the money that they ask for (Danquah and Longe 2011; Armstrong 2011). If they disobey the mallam, however, they will be cursed with bad luck (Graphic Online 2009).

These cybercriminals are typically young men aged 16 to 30. In many cases, they do not act alone but instead work for crime bosses who are thought to run both country-specific teams and broader regional ones across the West Africa region (Kavanagh 2013). For example, it appears that 419 scams have gone from national-level operations to regional syndicates, and some Nigerians engaged in 419 scams in their own country may have relocated some of their activities to Ghana as part of this expansion (ibid.).

An Evolution toward More Technically Sophisticated Attacks: Phishing and Mobile Banking Hacks

While cybercriminal groups in Ghana are still actively engaged in scamming, they are also adopting some of the new forms of cybercrime emerging around the world that make use of more technologically sophisticated techniques. There are currently a high number of phishing attacks in Ghana. These make up the second-largest category of reports to the e-Crime Bureau, accounting for more than 50 incidents. Phishing consists of cybercriminals sending emails made to appear as if they are from legitimate sources in order to trick victims into clicking on a link in the message. The links typically direct victims to fake websites that ask them to input passwords or credit card details, enabling the cybercriminals to steal confidential information on business networks, access individuals' bank accounts or use their credit cards, or other forms of theft.

Cybercriminals in Ghana are also increasingly making use of malware that targets smartphones, given that smartphones are the prime means of Internet access for many in the country. Security threats to mobile devices and malware attacks combined make up the third-largest category of reports to the e-Crime Bureau, accounting for more than 40 incidents in total. In particular, there has been a recent spate of incidents involving mobile banking. The mobile banking sector in Ghana has expanded considerably in recent years, with many of the country's major banks now offering some form of mobile banking. The potential to steal significant funds from banks, combined with the relative cyber security weaknesses of mobile devices, has thus made mobile banking a particularly attractive target for cybercriminals.

Other common attacks include website defacements and DDoS attacks to shut down websites. There are also a large number of botnets emanating in the region. These account

for a small proportion of reports to the e-Crime Bureau, however.³

Challenges for Combatting Cybercrime: Identifying the Root Causes of the Problem

Poverty and Unemployment

High poverty and unemployment levels in the country are driving some Ghanaians to engage in cybercrime. With close to one-third of young people unemployed in the country, Internet crime offers the possibility of earning large sums of money (Morton 2011). A distinctive sakawa culture, in which cybercriminals engage in lavish displays of their wealth, further entices them. Sakawa fraudsters typically wear flashy clothes and drive expensive cars. Observing this, other young men want to engage in cybercrime in order to enjoy this same lifestyle.

Engaging in cybercrime also enables those who feel powerless and live in poverty to enact retribution against the groups that they believe have been or are exploiting them. Some hold the West — and notably the legacy of colonialism and slavery — responsible for the ills that they suffer today. Others blame the government and wealthy Ghanaians, who are often corrupt. By targeting these groups, cybercriminals may feel that they are obtaining justice.

Electronic Waste

Actions by developed countries are also contributing to cybercrime in Ghana. The sending of electronic waste, which consists of old computers, monitors, cell phones and other electronic devices, to Ghana for disposal has provided cybercriminals with ready access to the data remaining on these devices — data that they can use to engage in cybercrime. To reduce the cost of recycling or disposing of electronic waste, companies in developed countries are increasingly shipping it to the developing world for handling. Ghana is one of the main recipients and receives some 215,000 tons of electronic waste each year, much of it consisting of computers and monitors (Amoyaw-Osei et al. 2011). The majority of these imports come from the United Kingdom (up to 60 percent), with the United States a close second (Doyon-Martin 2015).

Upon arrival in Ghana, the used computers and other electronics are sorted into categories: those that still work or can be repaired are reconditioned and then sold in markets for used goods. Those that cannot be are sent to dump sites. Agbogbloshie, located in a suburb of Accra,

3 Other types of crimes committed using the Internet in Ghana include use of social media for propaganda by terrorist groups, such as Boko Haram and Al Qaeda's use of the dark web for communication; human trafficking rings posting of false recruitment ads online to lure victims; and online child pornography. However, discussion of these activities is beyond the scope of this chapter, which focuses on cybercrime carried out for direct economic gain.

has become one of the world's largest dumping grounds for used computers and other electronic waste. At the dump sites, a secondary industry exists in which workers — typically children living in slums — sort through the electronic waste to extract copper, aluminum and other materials.⁴

This creates an opportunity for cybercriminals because the hard drives of these computers and other devices often contain credit card numbers, bank account information, business or personal documents, email addresses of colleagues and family and other confidential information belonging to the previous owners of the machines. The memories of mobile phones, too, often contain large amounts of data — particularly as they increase in computing power.

In many instances, obtaining data from these devices is made easier for cybercriminals because the data has not been wiped from the machines. This might be either because the previous owner did not erase the hard drive before disposing of the machine or else because the companies to whom they entrusted their devices for recycling assured them that they would clean them properly and then failed to do so. Moreover, even when a hard drive has been expunged, it is generally still possible to retrieve information from the hard drive if the hard drive has not been physically destroyed.

In one case, a media investigation found a computer in Ghana that had belonged to Northrop Grumman, one of the largest US military contractors. It contained confidential data about \$22 million in US government contracts (Klein 2009). Used machines formerly owned by the US Army, Homeland Security and other US government departments have also been found in Ghana.

The data on such machines has provided a boon for cybercriminals, who can purchase computers from second-hand vendors for as little as US\$27 to \$40 or comb through dump sites to find them (Stewart 2011). They then use the information they have found to target the original owners or their business associates and relatives. In one of the most well-publicized examples, a Ghanaian cybercriminal obtained the discarded hard drive of US Congressman Robert Wexler and threatened to sell his social security number to identity thieves if Wexler did not pay him (Warner 2011).

Insufficient Institutional Expertise and Funding

Institutions in Ghana — including law enforcement, the judiciary and government agencies — often have a limited understanding of cybercrime and thus lack the

expertise to effectively combat the problem. For example, the Commercial Crime Unit of the Ghana Police Service's Criminal Investigations Department (CID) is tasked with investigating and prosecuting cybercriminals. However, it often does not have the necessary technical skills to carry out digital forensics, which involve recovering and analyzing electronic evidence and preserving it in its most original form (Boateng et al. 2011).

Insufficient funding for law enforcement is a related challenge. This includes not enough funding for training police officers. Some recent training programs in digital forensics, cyber fraud detection and other cyber security-related skills have been beneficial, but more are needed. Lack of equipment is a key issue as well. For instance, the police force does not have an adequate computer lab in which to conduct digital forensics.

When law enforcement is able to find and hire technically skilled personnel, retaining them has proved difficult. Some officers, after having undergone a recent cyber security training program, left shortly afterwards to join the private sector, which offered them significantly higher salaries for their newly acquired skills.

An Insufficient and Unclear Legal Regime

The country's current laws are insufficient and often unclear with respect to cybercrime. This makes it difficult for police, prosecutors and judges to determine how the law applies to various offences. For example, the Economic and Organised Crime Office Act (Republic of Ghana 2010), which established a specialized government body to investigate and prosecute economic and organized crime, tasked the body with investigating a host of crimes including "prohibited cyber activity." However, the act does not specify which particular cyber offences this might include. As another example, the Electronic Transactions Act (Ghana Trade Portal 2008), which is intended to facilitate electronic communications, does define the specific cyber offences covered by the act — but it does not specify procedures to ensure the integrity and admissibility of electronic evidence. As a result, members of the police force, prosecutors and judges often have different and frequently conflicting interpretations of the law regarding cybercrime.

Limited Collaboration

The problem is compounded by limited collaboration between institutions that should be working together to address the challenge. For instance, the Electronic Transactions Act requires ISPs to provide law enforcement with technical assistance in response to a court order (such as a suspect's Internet Protocol address history), but ISPs rarely do so. And when business, academia or government run projects or initiatives to combat cybercrime, they seldom coordinate with one another.

⁴ This activity typically involves working in hazardous conditions. To extract the materials, workers often burn the electronic waste, releasing toxic chemicals in the process.

There have been attempts by foreign governments and agencies to help tackle the cybercrime challenge through training and other capacity-building initiatives, but most of these initiatives have failed to generate concrete results. One reason for this is that some capacity-building initiatives target only one specific group. For instance, one recent program provided training for prosecutors without providing complementary training for CID detectives, who investigate cybercrime cases before sending them to prosecutors for trial.

Enforcement Challenges, including Corruption

Even in instances where there is legal clarity, enforcement of cybercrime laws is often weak. The large number of victims located abroad makes prosecution more difficult (Darko 2015). Corruption, too, contributes to the lack of enforcement and is also fuelling the cybercrime challenge. For instance, fraudsters may sometimes bribe law enforcement or judges to overlook their activities. Furthermore, the government has an incentive to ignore cybercrime since it provides Ghanaians with a source of income (Morton 2011).

Potential Solutions

Job Training and Development Programs for Youth

Given the number of unemployed young men turning to cyber fraud, there is a need for development programs targeted at youth. As evidenced by their success in committing online crime, fraudsters do have some technical skills, so such schemes could foster and utilize this information technology (IT) talent (Boateng et al. 2011). One possibility might be for the government to sponsor training programs for youth geared at preparing them to work in online outsourced roles for international companies.

Destroying or Wiping Data on Devices Before They Are Exported as Electronic Waste

To tackle cybercrime stemming from electronic waste, there must be greater efforts to destroy or wipe clean the memories of devices before they are exported. The best way to do this is to physically destroy the hard drives with a hammer or other blunt instrument. Software that erases the hard disks can also be effective. The first step is for governments of countries that export electronic waste to launch a public education campaign. Many people are unaware that the data remaining on their used devices could end up in the hands of cybercriminals. They may also not know how to safely destroy or wipe the information on these devices. It is thus essential to encourage individual responsibility in the proper disposal of used devices and to share the knowledge of how to do so.

Companies also need to behave responsibly regarding the safe disposal of used devices. The second step is therefore for governments to sensitize recycling firms that export electronic waste to the associated cybercrime dangers; they must be made more aware of the importance of destroying or wiping all hard drives before shipping them. If necessary, governments may need to require them to do so.

Further Restricting the Export of Electronic Waste

Firms — not just recycling firms that export electronic waste but those that manufacture devices — should also be encouraged to recycle used devices instead of exporting them, as this is the most effective way to ensure that cybercriminals cannot obtain and exploit them. Recycling firms lack economic incentives to do this, however. It is significantly cheaper for them to send the devices intact to the developing world as electronic waste rather than go through the costly process of disassembling the devices, disposing of toxic substances and recuperating the materials that can be reused.

Public opinion can play an important role in incentivizing manufacturers, however. Apple, extremely conscious of its brand, has developed one of the largest recycling programs for used devices in the industry in order to bolster its “green” image. Governments can mandate that companies recycle a certain percentage of used devices or, if they already do require it, can increase the amount required. In the United States, for instance, most states have legislated that device manufacturers must pay the cost of recycling part of their electronic waste, but that percentage is often too low (Risen 2016).

Tightening regulations on the export of electronic waste will be essential too. Although many countries ostensibly ban electronic waste exports, in practice this is often flouted: the 1992 Basel Convention on the Control of Transboundary Movements of Hazardous Wastes and their Disposal, which 182 countries are party to, prohibits developed countries from exporting hazardous waste, including electronic waste, to developing countries in most instances. Many companies, however, skirt national regulations by classifying electronic waste as “donations of second-hand goods.” Governments must close this loophole if the Convention is to be effective.

Moreover, the biggest exporter of electronic waste to Ghana — the United States — is not bound by the Convention. Although the United States signed the agreement, ratification has been gridlocked in Congress. It is the only industrialized country in the world not to have ratified the Convention (ibid.). The United States needs to do so to stem the flow of electronic waste to Ghana and the cybercriminal activity that derives from it.

Regulatory Measures

To address the insufficiency of current legislation regarding cybercrime, Ghana should consider incorporating new measures into law such as the legal right to asset seizure of the proceeds of cybercrime, either by introducing new legislation or by amending existing legislation. Doing so would not only give law enforcement and judges greater power to punish known cybercriminals but also serve as a deterrent to others.

In order to tackle the lack of clarity of current laws involving cybercrime, Ghana should also sign the Convention on Cybercrime (the Budapest Convention), an international treaty on Internet and computer crime that harmonizes national laws. This would align Ghanaian law with international norms and improve the coherence of the country's legislation. In June 2016, Ghana was invited to accede to the Budapest Convention; doing so is essential if the country is to effectively curb cybercrime.

Increasing Knowledge and Understanding of Cybercrime among All Stakeholders, including Bolstering Technical Skills and Funding for Law Enforcement

Increasing knowledge and understanding of cybercrime among all stakeholders is essential. Given that law enforcement personnel often have insufficient technical knowledge, more training programs in digital forensics and other related areas are needed. It is also essential to recruit professional IT experts when possible.

To achieve this, the government must allocate greater funding to the CID. This, in turn, will require persuading the government to make the fight against cybercrime a key priority. The government has begun to pay increased attention to the issue since some international companies restricted the use of credit cards in Ghana because of high incidences of cyber fraud. As a result, the government is beginning to realize that cybercrime is a problem for the country's burgeoning e-commerce industry and its international business reputation, and thus for the economy as a whole. Despite this, greater efforts are needed to increase the government's awareness of the cybercrime threat. One way to do so would be to conduct an economic study to attempt to quantify the business revenue lost to Ghana because of cybercrime.

Engaging All Stakeholders

Effectively tackling cybercrime will require the cooperation of all stakeholders: police, prosecutors, the judiciary, government ministries and agencies, Parliament, the private sector (including ISPs) and civil society. The government has made important progress recently: the Ministry of Communications launched the country's National Cyber Security Policy and Strategy in 2015 and held a validation workshop with a range

of stakeholders, including ISPs. With the help of the International Telecommunication Union, Ghana also set up a Computer Emergency Response Team in 2014 to coordinate cyber security incident response. However, much more remains to be done.

There is a need for greater collaboration between all stakeholders, including for improved information sharing between ISPs and law enforcement as well as between law enforcement and the judiciary. Steps to achieve this might involve either the government or civil society convening regular fora on cybercrime that bring together all stakeholders to discuss the challenges.

An additional method of enhancing coordination between law enforcement and the judiciary could involve developing a best practices document on the handling of electronic evidence. Such a document is needed in order to standardize the process of handling cybercrime evidence across the country and to ensure the integrity of electronic evidence in cybercrime investigations and prosecutions.

International cooperation is key as well. Given that cybercrime transcends national borders, the Ghanaian police needs to work more closely with members of ECOWAS as well as with the broader international community in order to more effectively conduct cross-border investigations and evidence collection.

A POSSIBLE OVERARCHING STRATEGY

The closely linked nature of Internet infrastructure development and cybercrime means that policy makers may want to consider deploying an overarching strategy that encompasses both. In particular, they should consider the following measures.

Leveraging Multiplier Effects

This analysis has pointed to certain root causes that are common to both Internet infrastructure development challenges and the cybercrime problem. Specifically, poverty and corruption feature as root causes of both. This suggests that, when determining where to invest limited resources, one beneficial strategy for policy makers may be to concentrate their efforts on some of the shared root causes of both problems because this will have a multiplier effect. Thus, investing in programs that target poverty alleviation or anticorruption — which tackle shared causes of both issues — are likely to have a larger impact on the country's well-being than programs that focus on independent causes.

A Joint Approach

Taking a joint approach to solving these problems in discussions and workshops could prove highly beneficial. For example, this analysis has also pointed to a solution that is common to both Internet infrastructure development challenges and cybercrime problems: that is, the need to bring all stakeholders in each case together. Not only are the problems interlinked, but there is also significant overlap in terms of the major stakeholders who are involved. In addition to regular meetings between stakeholders for each issue suggested earlier in this chapter, it might also be beneficial to hold a number of joint meetings so as to approach these issues in a holistic manner.

Joint meetings would generate ideas on how best to tackle these challenges together in order to be most effective. For example, if a lack of legitimate jobs is a significant factor contributing to cybercrime and if developing Internet infrastructure is key to the country's economic growth, then it might make sense to launch a public works program that employs people to work on large Internet infrastructure projects. In this way, it would be possible to provide employment that would help reduce cybercrime and simultaneously improve the country's Internet infrastructure, which would in turn stimulate economic growth.

CONCLUSIONS

Further promoting the development of Internet infrastructure in Ghana is central to the country's continued economic growth and prosperity. Yet increases in Internet connectivity will — if corresponding measures are not employed to keep cybercrime at bay — also result in an increase in cybercrime. This chapter suggests that an overarching strategy that combines leveraging multiplier effects (i.e., concentrating efforts on some of the shared root causes of both problems) and a joint approach (i.e., holding a number of joint meetings and workshops to approach these interlinked issues in a holistic manner) would be the most effective means of improving well-being.

Some areas for further study include additional research into common root causes — beyond poverty and corruption — of both Internet infrastructure development challenges and cybercrime problems. Supplementary work is also needed to consider potential joint solutions as well. Joint meetings bringing together all major stakeholders would be an ideal venue in which to do this sort of brainstorming.

Although this chapter focuses on Ghana, many of the findings can likely be applied to neighbouring countries in the West Africa region — including Nigeria and Cameroon, which have markedly similar characteristics. In some cases, they can be extrapolated for developing

nations more broadly, given that a number of emerging countries are facing similar issues. Given that cybercrime impacts all countries around the world, irrespective of their level of development, it is clear that addressing the dual challenges of Internet infrastructure development and cybercrime is an urgent priority.

ACKNOWLEDGEMENT

The authors are grateful to the UK Foreign and Commonwealth Office for funding this work.

WORKS CITED

- Abbey, Emelia Ennin. 2016. "Vodafone bemoans rising cable theft." *Graphic Online*, April 26. www.graphic.com.gh/business/business-news/vodafone-bemoans-rising-cable-theft.html.
- Abdul-Jalil, Yakubu. 2015. "Google Ghana Launches Metro Fibre Project." *Graphic Online*, October 3. <http://graphic.com.gh/news/general-news/50586-google-ghana-launches-metro-fibre-project.html>.
- Abubakar, Zulaihatu. 2012. "Sakawa Guy Confesses." *Modern Ghana*, September 22. www.modernghana.com/news/419261/1/sakawa-guy-confesses.html.
- Acquaye, Nana Appiah. 2013. "The Rise and Fall of Ghana's Biggest Internet Café." *BiztechAfrica*, February 14. www.biztechafrica.com/article/rise-and-fall-ghanas-biggest-Internet-cafe/5320/#.VoH31FK_Gbg.
- . 2014. "Ghana Roads Ministry: Cable Cuts a Concern." *BiztechAfrica*, February 4. www.biztechafrica.com/article/ghana-roads-ministry-cable-cuts-concern/7647/#.VoCPZVK_Gbg.
- . 2015. "NCA to Auction Infrastructure License." *BiztechAfrica*, August 7. www.biztechafrica.com/article/nca-auction-infrastructure-license/10434/#.VoL9UVK_Gbg.
- Adam, Basiru. 2014. "What Fibre Cuts Do to You." *B&FT Online*, July 30. <http://thebftonline.com/business/ict/11896/what-fibre-cuts-do-to-you.html>.
- Amoyaw-Osei, Y., O. O. Agyekum, J. A. Pwamang, E. Mueller, R. Fasko and M. Schlupe. 2011. *Ghana e-Waste Country Assessment: SBC e-Waste Africa Project*. Coordinated by the Basel Convention. March.
- Ampomah, Eunice Hilda. 2016. "Vodafone sets up cable theft courts." *Ghana News Agency*, April 20. www.ghananewsagency.org/social/vodafone-sets-up-cable-theft-courts-102918.
- Armstrong, Alice. 2011. "Sakawa Rumours: Occult Internet Fraud and Ghanaian Identity." Working Paper No. 8/2011. University College London, Department of Anthropology. www.ucl.ac.uk/anthropology/research/working-papers/082011.pdf.
- BiztechAfrica. 2013. "Ghana Telecoms, Roads Collaborate to Minimise Cable Cuts." March 13. www.biztechafrica.com/article/ghana-telecoms-roads-collaborate-minimise-cable-cu/5659/#.VRgIvGbZfAo.
- Boakye, Kojo. 2014. "Affordable Internet in Ghana: The Status Quo and the Path Ahead." A4AI (Alliance for Affordable Internet). https://a4ai.org/wp-content/uploads/2014/07/Ghana-Case-Study_FINAL.pdf.
- Boateng, Richard, Longe Olumide, Robert Stephen Isabalija and Joseph Budu. 2011. "Sakawa — Cybercrime and Criminality in Ghana." *Journal of Information Technology Impact* 11 (2): 85–100. www.jiti.com/v11/jiti.v11n2.085-100.pdf.
- Calandro, Enrico, Christoph Stork and Alison Gillwald. 2012. "Internet Going Mobile: Internet Access and Usage in 11 African Countries." Policy Brief No. 2. Research ICT Africa. www.researchictafrica.net/publications/Country_Specific_Policy_Briefs/Internet_going_mobile_-_Internet_access_and_usage_in_11_African_countries.pdf.
- Citifmonline. 2016a. "Govt backtracks on withdrawal of 20% tax on mobile phones." *Citifmonline.com*, February 4. <http://citifmonline.com/2016/02/04/govt-draws-back-on-20-tax-on-mobile-phones/>.
- . 2016b. "Telecom analyst demands tax cuts on imported phones." *Citifmonline.com*, August 11. <http://citifmonline.com/2016/08/11/telecom-analyst-demands-tax-cuts-on-imported-phones/>.
- CRU Wire and Cable News. 2014. "Ghana Solves Theft Problems with Overhead Fibre Optic Cable," November 18. <http://wireandcablenews.crugroup.com/wireandcablenews/news/free/2014/11/2071615/>.
- Danquah, Paul and O. B. Longe. 2011. "Cyber Deception and Theft: An Ethnographic Study on Cyber Criminality from a Ghanaian Perspective." *Journal of Information Technology Impact* 11 (3): 169–82. www.jiti.net/v11/jiti.v11n3.169-182.pdf.
- Darko, Sammy. 2015. "Inside the World of Ghana's Internet Fraudsters." *BBC Africa*, May 10. www.bbc.com/news/world-africa-32583161.
- Dela Klutse, Felix. 2015. "7.5m Ghanaians live on GH¢3 daily." *JOY Online*, March 16. www.myjoyonline.com/business/2015/march-16th/75m-ghanaians-live-on-gh3-daily.php.
- Doyon-Martin, Jacquelynn. 2015. "Cybercrime in West Africa as a Result of Transboundary E-Waste." *Journal of Applied Security Research* 10 (2): 207–20. www.tandfonline.com/doi/pdf/10.1080/19361610.2015.1004511.
- Foster, William, Seymour Goodman, Eric Osiakwan and Adam Bernstein. 2004. "Global Diffusion of the Internet IV: The Internet in Ghana." *Communications of the Association for Information Systems* 13: 654–81.

- Frempong, Godfred. 2012. "Understanding What Is Happening in ICT in Ghana: A Supply- and Demand-side Analysis of the ICT Sector." Evidence for ICT Policy Action Policy Paper 4, Research ICT Africa. http://researchictafrica.net/publications/Evidence_for_ICT_Policy_Action/Policy_Paper_4_-_Understanding_what_is_happening_in_ICT_in_Ghana.pdf.
- Ghana Business News. 2015. "Ghana Road Fund Carries forward over GH¢230m Indebtedness — Minister." March 28. www.ghanabusinessnews.com/2015/03/28/ghana-road-fund-carries-forward-over-gh%2%A2230m-indebtedness-minister/.
- Ghana News Agency. 2014. "MTN Worried over Fibre Optic Cuts." July 4. www.ghananewsagency.org/science/mtn-worried-over-fibre-optic-cuts--76839.
- . 2016. "Trade Ministry empowers Vodafone to fight cable theft." Modern Ghana, August 30. www.modernghana.com/news/716068/trade-ministry-empowers-vodafone-to-fight-cable-theft.html.
- Ghana Trade Portal. 2008. *Electronic Transactions Act (Act 772), 2008*. www.ghanatrade.gov.gh/Laws/electronic-transactions-act-act-7722008.html.
- GhanaWeb. 2012a. "Community Members Act to Stop Theft of Vodafone Cables." October 31. www.ghanaweb.com/GhanaHomePage/regional/artikel.php?ID=254912.
- . 2012b. "Vodafone Cable Thieves Arrested by Ghana Police." October 10. www.ghanaweb.com/GhanaHomePage/NewsArchive/artikel.php?ID=252728.
- . 2015a. "Four Vodafone Staff Jailed over Cable Theft." May 19. www.ghanaweb.com/GhanaHomePage/crime/4-Vodafone-staff-jailed-over-cable-theft-358740.
- . 2015b. "Fibre-optic Operators Call for Cable Protection Law." November 30. www.ghanaweb.com/GhanaHomePage/NewsArchive/Fibre-optic-operators-call-for-cable-protection-law-397325.
- . 2016. "Vodafone cracks down on cable theft." April 25. www.ghanaweb.com/GhanaHomePage/crime/Vodafone-cracks-down-on-cable-theft-433623.
- Graphic Online. 2009. "Sakawa rituals paralyse student." Modern Ghana, March 27. www.modernghana.com/news/208469/1/sakawa-rituals-paralyse-student.html.
- . 2012. "Telcos Co-Share Ducts to Lay Fibre Cables." Modern Ghana, August 28. www.modernghana.com/news/414121/1/telcos-co-share-ducts-to-lay-fibre-cables.html.
- IT News Africa. 2015. "Ghana: Alcatel-Lucent Completes Fibre Backbone Project." May 20. www.itnewsafrika.com/2015/05/ghana-alcatel-lucent-completes-fibre-backbone-project/.
- JOY Online. 2013. "Government to Ban Copper Export — Minister." Modern Ghana, June 4. www.modernghana.com/news/467023/1/government-to-ban-copper-export-minister.html.
- . 2014. "MTN Delivers Digital Advantage to Customers through Aggressive Network Investments." May 8. www.myjoyonline.com/business/2014/May-8th/mtn-delivers-digital-advantage-to-customers-through-aggressive-network-investments.php.
- Kavanagh, Camino, ed. 2013. "Getting Smart and Scaling Up: Responding to the Impact of Organized Crime on Governance in Developing Countries." New York University, Center for International Cooperation. Report, June. http://cic.nyu.edu/sites/default/files/kavanagh_crime_developing_countries_report_w_annexes.pdf.
- Klein, Peter. 2009. "Ghana: Digital Dumping Ground." FrontlineWorld. www.pbs.org/frontlineworld/stories/ghana804/video/video_index.html.
- Kunateh, Masahudu Ankiilu. 2015. "Ghana: Telcos Welcome New Directive to Halt Fibre Cuts." *All Africa/The Chronicle*, January 7. <http://allafrica.com/stories/201501071449.html>.
- Morton, Thomas. 2011. "The Sakawa Boys: Inside the Bizarre Criminal World of Ghana's Cyber-Juju Email Scam Gangs." Motherboard, April 5. <http://motherboard.vice.com/read/the-sakawa-boys-inside-the-bizarre-criminal-world-of-ghanas-cyber-juju-email-scam-gangs>.
- Mustapha, Suleiman. 2014. "Engage Telcos in Road Construction." Graphic Online, December 15. <http://graphic.com.gh/news/general-news/35454-engage-telcos-in-road-construction.html>.
- Naphtal, Akin. 2015. "Telcos in Ghana Welcome New Directive to Minimise Fibre Cuts." *Mobile World Mag*, January 8. <http://mobileworldmag.com/telcos-ghana-welcome-new-directive-minimise-fibre-cuts/>.
- National Communications Authority. 2015. "News Item: Mobile Data Figures for the Month of July 2015." www.nca.org.gh/downloads/Data_Market_Figures_July2015.pdf.

- National Crime Agency. 2014. "Romance Fraud Mastermind Jailed in Ghana." National Crime Agency, October 31. www.nationalcrimeagency.gov.uk/news/478-romance-fraud-mastermind-jailed-in-ghana.
- Ogundeji, Olusegun Abolaji. 2014. "Ghana's Import Tax Removal on Smartphones Expected to Boost Local Africa Production." PCWorld.com, December 10. www.pcworld.com/article/2858172/ghanas-import-tax-removal-on-smartphones-expected-to-boost-local-africa-production.html.
- Republic of Ghana. 2010. *Economic and Organised Crime Office Act (Act 804), 2010*. <http://fic.gov.gh/wp-content/uploads/2015/11/EOCO-Act-804.pdf>.
- Risen, Tom. 2016. "America's Toxic Electronic Waste Trade." US News and World Report, April 22. www.usnews.com/news/articles/2016-04-22/the-rising-cost-of-recycling-not-exporting-electronic-waste.
- Sikiti da Silva, Issa. 2014. "Ghana Govt Worried about Rising Cybercrime." BiztechAfrica, June 5. www.biztechafrica.com/article/ghana-govt-worried-about-rising-cybercrime/8250/#.V4fuUhiG6T8.
- Stewart, Samantha. 2011. "Ghana's e-Waste Dump Seeps Poison." Newsweek.com, July 25. <http://europe.newsweek.com/ghanas-e-waste-dump-seeps-poison-68385?rm=eu>.
- Warner, Jason. 2011. "Understanding Cyber-Crime in Ghana: A View from Below." *International Journal of Cyber Criminology* 5 (1): 736–49.

ABOUT THE AUTHORS

Caroline Baylon served as the lead researcher on cyber security at Chatham House in London, United Kingdom, from 2013 to 2015, and was also editor of the institute's *Journal of Cyber Policy*, a peer-reviewed academic journal published by Routledge, Taylor & Francis. Caroline recently worked as an independent contractor carrying out research projects on cyber security for the UK Foreign and Commonwealth Office, looking at cyber proxy actors and at limiting cyber weapons proliferation. She is currently the information security research lead within the research-and-development section at AXA in Paris, France, and London, United Kingdom, which is establishing an internal think tank on cyber security. Her work there includes a research stream on cyber security issues impacting Africa. She has also worked as an independent consultant for a number of intelligence providers on cybercrime issues involving Sub-Saharan Africa. Caroline holds an M.Sc. in social science of the Internet from Balliol College, University of Oxford, and a B.A. in economics from Stanford University.

Albert Antwi-Boasiako is the founder of e-Crime Bureau, a cyber security firm based in Ghana, and a cyber security expert with the Interpol Global Cybercrime Expert Group. A graduate of the University of Trento, Italy, and the University of Portsmouth, United Kingdom, Albert is currently a Ph.D. research fellow with the University of Pretoria, South Africa. Albert has conducted different cyber security capacity-building projects for a number of international organizations including the Council of Europe, United Nations Office on Drugs and Crime, United Nations Conference on Trade and Development, Commonwealth Cybercrime Initiative and the Inter-Governmental Action Group against Money Laundering in West Africa. Albert is a visiting lecturer in cybercrime, cyberterrorism and cyber security at the Kofi Annan International Peacekeeping Training Centre, Accra, Ghana, and a research associate with the African Centre for Cyberlaw and Cybercrime Prevention based in Kampala, Uganda.

**CHAPTER SEVEN:
CRITICAL INFRASTRUCTURE AND THE INTERNET OF THINGS**
Tobby Simon

Copyright © 2017 by Tobby Simon

ACRONYMS

CERTs	Computer Emergency Response Teams
CIS	critical infrastructure systems
DDoS	distributed denial of service
DoS	denial of service
ICS	industrial control system
IIoT	Industrial Internet of Things
IoT	Internet of Things
IP	Internet Protocol
IT	information technology
SCADA	supervisory control and data acquisition

INTRODUCTION

In November 2015, US prosecutors indicted three men in connection to the massive 2014 JPMorgan Chase cyber attack and the hacking of several other financial institutions. The vast, multi-year criminal enterprise centred on compromised private information involving 100 million institutional customers, which fuelled a web of stock manipulation, credit-card fraud and illegal online gambling. The globe-trotting conspiracy hacked servers in various countries, and in one instance exploited the notorious Heartbleed bug. With the stolen data, the group defrauded investors by criminally manipulating stocks, artificially inflating them. They deceived private companies into offering their shares publicly. The group then carefully manipulated the stock prices of the publicly traded companies, spammed email “tips” to institutional clients using stolen information, then quickly would sell off for profit, causing the stock values of the companies they had misled to collapse. The group illegitimately earned millions of dollars in this manner (Farrell and Hurtado 2015).

In a case study paper, Robert M. Lee, Michael J. Assante and Tim Conway (2014) provide an account of a cyber attack on a German steel mill:

In December, 2014 the German government’s Bundesamt für Sicherheit in der Informationstechnik (BSI) (translated as Federal Office for Information Security) released their annual findings report. In one case they noted that a malicious actor had infiltrated a steel facility. The adversary used a spear phishing email to gain access to the corporate network and then moved into the plant network.

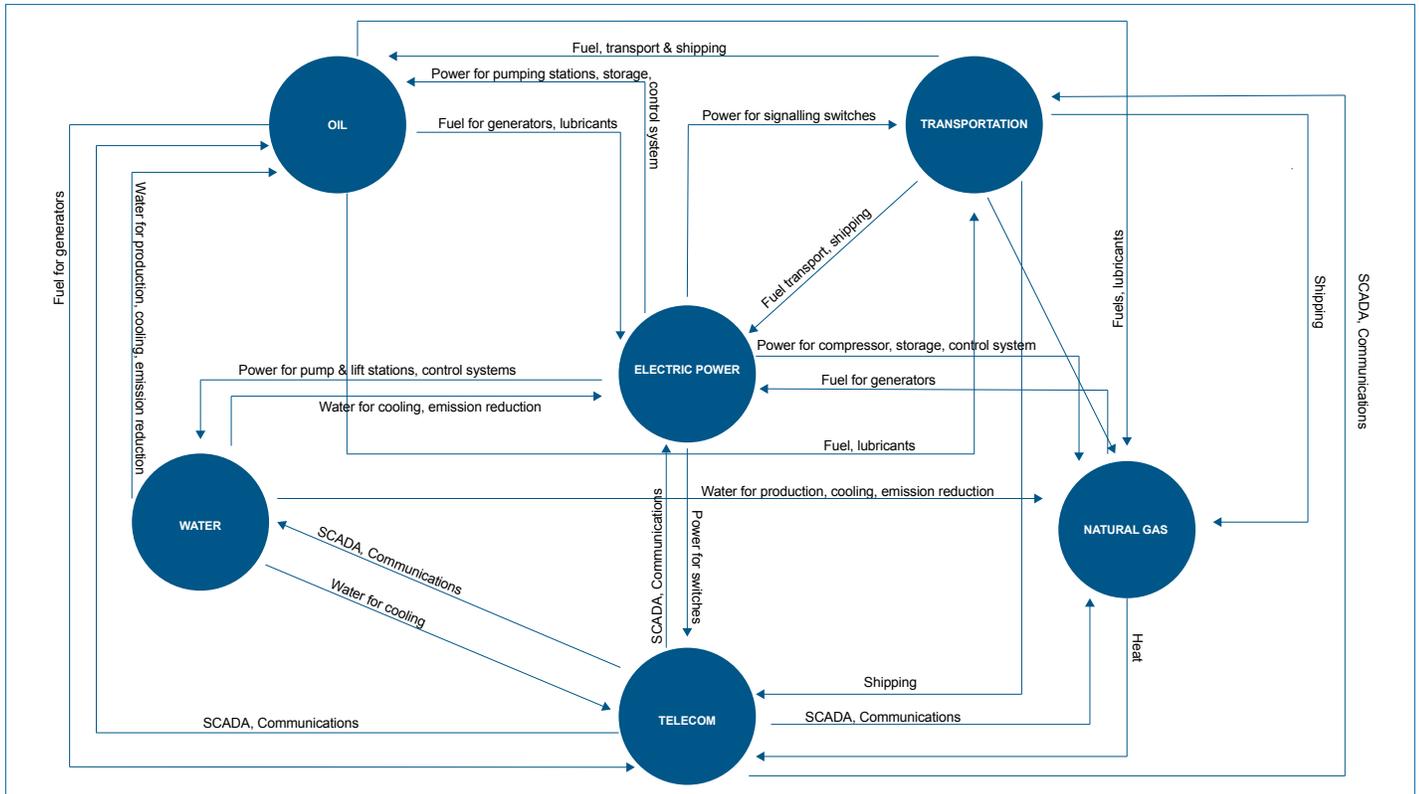
According to the report, the adversary showed advanced knowledge of ICS [industrial control system] and was able to cause multiple components of the system to fail. This specifically impacted critical process components to become unregulated, which resulted in massive physical damage.

There have been other cases where hackers have used printers, thermostats and videoconferencing equipment to breach security systems. Cybercrime costs the global economy some CDN\$400 billion per annum (Desjardins 2015). In recent years, cyber attacks on Sony, the retailer Target and the Internet dating site Ashley Madison have shown that the technology that offers so many opportunities also brings with it significant threats. Data breaches are usually not identified immediately, as seen in the JPMorgan Chase case, where it was only much later determined that hacked contact information was used in stock manipulation. While details about the damage caused by the attack on the German steel facility are not known, the incident leads to speculation regarding the prospective impact of a larger, more organized cyber attack on the nation’s critical infrastructure.

Internet-enabled infrastructure has transformed the boundaries of Internet technology, be it through home-automation concepts, energy-management systems and “smart homes”; wellness devices and network-enabled medical gadgets, which are revolutionizing health care sectors; intelligent vehicles, networked traffic systems and road and bridge sensors; or innovations in agricultural, industrial and energy production and distribution. The rise of “smart cities” has been increasing access to and the availability of information manifold. However, while this has opened up myriad avenues for efficiency, and is helping reap benefits to the tune of billions of dollars for the global economy, the unfettered rise of the Internet of Things (IoT) raises a plethora of issues: the IoT brings with it a concomitant set of concerns about the security and privacy of people, telecoms networks and power utilities, say, through illegitimate breaches of the networks undergirding critical infrastructure, as the efficiency of Internet connectivity also accelerates susceptibility to security violations through the misuse of IoT data. A “promise vs. peril” discussion has subsequently emerged within governmental and academic debates, which have begun to seek the best means to address the complex interdependence between critical infrastructure and IoT systems.

CRITICAL INFRASTRUCTURE

The term “infrastructure” is evolutionary and is often ambiguous. It is traditionally defined as any physical asset that is capable of being used to produce services or support

Figure 1: Example of Infrastructure Interdependencies

Source: Author.

the structure and operation of a society or an enterprise. Today, the notion of public infrastructure has broadened and encompasses such structures as roadways, bridges, airports and airway facilities, mass transportation systems, waste treatment plants, energy facilities, hospitals, public buildings and space or communication facilities, for example (Moteff and Parfomak 2004).

Critical infrastructure, on the other hand, includes physical and virtual facilities and services that form the basis for a nation's defence, a strong economy and the health and safety of its citizens. It is important as it provides necessities such as water and food, electricity and gas, telecommunications and broadcasting, health services, the financial system and the transportation system. They are essential for social cohesion and economic performance (see Figure 1).

At the heart of critical infrastructure is an ICS, which includes supervisory control and data acquisition (SCADA) systems, and other types of control systems that monitor processes and control flows of information. The functionality of an ICS is like the on or off feature of a light switch. For instance, an ICS can regulate the flow of natural gas to a power generation facility or the flow of electricity from a grid to a home.

An ICS is a proprietary and — most often — closed system. As an isolated, so-called air-gapped system, it is

not vulnerable to virtual attacks, although it is susceptible to attacks by way of physical access, such as from infected removable devices (for instance, if an employee or supplier unwittingly uses an infected USB device within an air-gapped system). As technology continues to grow, more ICSs are connected to the Internet. This makes them vulnerable to multifarious attacks.

The operating environment for critical infrastructure is increasingly complex, driven by a number of factors, including globalization, the evolution of technology and the interconnected nature of critical infrastructure supply chains, networks and systems. This complexity, in particular, impacts the ability to understand and manage cross-sector dependencies (Brandis 2014).

Computers and communications, themselves critical infrastructures, are increasingly tying other infrastructures together. The growing interconnectedness from networking means that a disruption in one network may lead to disruption in another. This reliance on computers and networks increases critical infrastructure's vulnerability to cyber attacks (Moteff and Parfomak 2004).

TYPES OF CYBER ATTACKS

Cyber attacks can be divided into four main groups — "hacktivism," cybercrime, cyber espionage and cyberwar,

although the lines often blur — i.e., hackers may also engage in cybercrime or cyber espionage. Moreover, what may be considered hacktivism in one nation could be considered intelligence or cyberwar in another nation. It should, therefore, be noted that the categories intersect with each other despite theoretical delineation.

Hacktivism

Hacktivism emerged in the late 1980s in the form of Internet viruses and worms spreading political propaganda and messages of protest. The group Worms Against Nuclear Killers is an example of early hacktivism; in 1989, these Australia-based anti-nuclear hackers installed worms into the networks of NASA and the US Department of Energy to protest the launch of a space shuttle carrying radioactive plutonium. By the middle of the next decade, denial-of-service (DoS) attacks became common, often taking the form of message or traffic floods; for example, in 1994, the “Zippies” group spammed email accounts in the United Kingdom to protest against a bill that outlawed outdoor dance music festivals. Dorothy Denning (2015) writes that the term “hacktivism” was coined in 1996 by the Cult of the Dead Cow hackers’ group, and the term picked up media momentum during the 1998-1999 Kosovo conflict, when DoS attacks were launched against websites in those member countries participating in the North Atlantic Treaty Organization’s aerial bombardment of Yugoslavia. Hacktivism has become a common means of protest: groups exist worldwide, some associating themselves with a specific country, such as Anonymous Syria, others associate themselves with a particular government or a political group, such as Cyber Caliphate, while others express no particular allegiance, such as Anonymous. Anonymous, a loosely organized group of hacker activists known for wearing Guy Fawkes masks, garnered popularity with the launch of Project Chanology, protests launched against what the group said was Internet censorship by the Church of Scientology. The group has since been responsible for cyber attacks and hacktivism against governments, terrorist organizations (including the Islamic State of Iraq and the Levant), corporations, religious groups and suspected sexual offenders, among others. Hacktivists, in addition to DoS attacks and defacing websites, often commandeer Twitter and Facebook accounts, make extensive use of social media to promote their actions and rally support, and steal and reveal sensitive information from the systems they penetrate (ibid.).

Cybercrime

Criminal hackers (motivated by economic gains through illegal penetration of computer networks, and relatively non-violent in nature) operate across the globe, replacing traditional forms of crime, costing the global economy an estimated CDN\$445 billion annually (Morag 2014).

Broadly, cybercrime includes fraud, sale in contraband and counterfeit items and online scams. Fighting cybercrime is particularly tricky because the crimes often challenge jurisdictional boundaries. A criminal hacker may sit in one country, use a server hosted in another and hack into systems housed in a third, rendering the legal and geographical components of the crime a challenge to investigate, let alone prosecute.

Cyber Espionage

Cyber espionage is a strategy aimed at obtaining critical governmental or corporate information by breaking into computer networks and systems. The strategy can be used to spy on any entity or group; for example, it is used for state-level purposes to understand rival country capabilities and attain classified information, or, in the case of industrial espionage, to gain access to rival business strategies and intellectual property. Cracking techniques and malicious software, such as a Trojan horse program, are employed to acquire personal, economic, military or political information through the Internet, computer networks or individual computers. Importantly, governmental or private actors sometimes undertake this even in the absence of hostilities. China has been particularly active in state-based hacking. According to one study, nearly half of all cyber-espionage attacks in the world originate from East Asia, in particular from China and North Korea (ibid.). North Korea, as mentioned in an example below, has waged distributed denial of service (DDoS) attacks on South Korea in the past decade. Also, Iran was blamed in 2013 for attacking Aramco, Saudi Arabia’s oil company, by erasing data from roughly 30,000 computers and penetrating Royal Saudi Navy and Marine Corps networks. Such operations are typically illegal in the victim entity, but may be launched or supported by a foreign state or an entity from abroad.

Cyberwarfare

Cyberwarfare has been defined as “actions by a nation-state to penetrate another nation’s computers or networks for the purpose of causing damage or disruption” (Clarke and Knake 2010), although the taxonomy has been widened to include non-state actors such as extremist groups, private firms, transnational criminal/terror groups and others. Countries are increasingly investing heavily in cyberwarfare technology, if not making cyber espionage a central aspect of their overall military strategy. This full-fledged threat to critical infrastructure is considered to have catapulted into a present danger with the discovery of the Stuxnet worm/virus in June 2010, and refers to any coordinated attacks waged against the critical infrastructure or control systems of a nation. Clandestine US attacks on the computer systems of Iranian nuclear enrichment facilities in 2012, as well as Russia’s cyber attacks against the websites and network infrastructures

Table 1: Threat Categories versus Impacts

Threat Type	Impact Type
Hacktivism	The interruption of life-sustaining services (minor)
Cybercrime	Economic damages (minor)
Cyber espionage	Economic damages (major) Severe degradation of national security
Cyberwar	The interruption of life-sustaining services (major) Economic damages (intermediate)

Source: Edwards (2004).

of Estonia and Georgia, are classified as tactics of cyberwar (Edwards 2004).

Table 1 outlines the impacts and relative severity of the four categories of threats. Where critical infrastructure is concerned, cyber espionage and cyberwar are far more harmful than hacktivism or cybercrime attacks, although they are perceived to be far less frequent (Morag 2014).

THE IoT

Walt Mossberg (2014) has described the IoT as a “constellation of inanimate objects [that] is being designed with built-in wireless connectivity, so that they can be monitored, controlled and linked over the Internet” (cited in Cha 2015). The IoT “refers to the connection of everyday objects to the Internet and to one another, with the goal being to provide users with smarter, more efficient experiences” (Cha 2015).

The Internet revolution has redefined the modern landscape and introduced unprecedented opportunity. The IoT has heralded “smart” living and is transforming every aspect of modern living, industry and the economy. Internet connectivity is now being built into a wide range of non-computer products, including kitchen and home appliances, lighting and heating products and insurance company-issued car-monitoring devices. These products contain three important components: an Internet connection, either in the device itself or in a base station; a digital sensor, to collect incoming data; and a processor, like any computing device. However as IoT industry develops, the threat landscape also changes drastically, augmenting information technology (IT) security concerns.

While the consumer IoT is set to revolutionize living, it comes with numerous risks. Recently, researchers from Proofpoint, a next-generation cyber security company,

reported that more than 100,000 smart TVs, refrigerators and other consumer items were compromised by hackers to transmit 750,000 malicious emails in a two-week period. Smart appliances are attractive to cybercriminals due to their 24-hour connectivity to the Internet and their poorly protected Internet environments (Prince Trust of India 2014). Researchers have shown how brakes in automobiles with on-board diagnostics, and other critical vehicular control systems, can be remotely controlled by virtually anyone with an Internet connection. One could take control of a such a vehicle by sending data to its interconnected entertainment and navigation system via a mobile phone network.

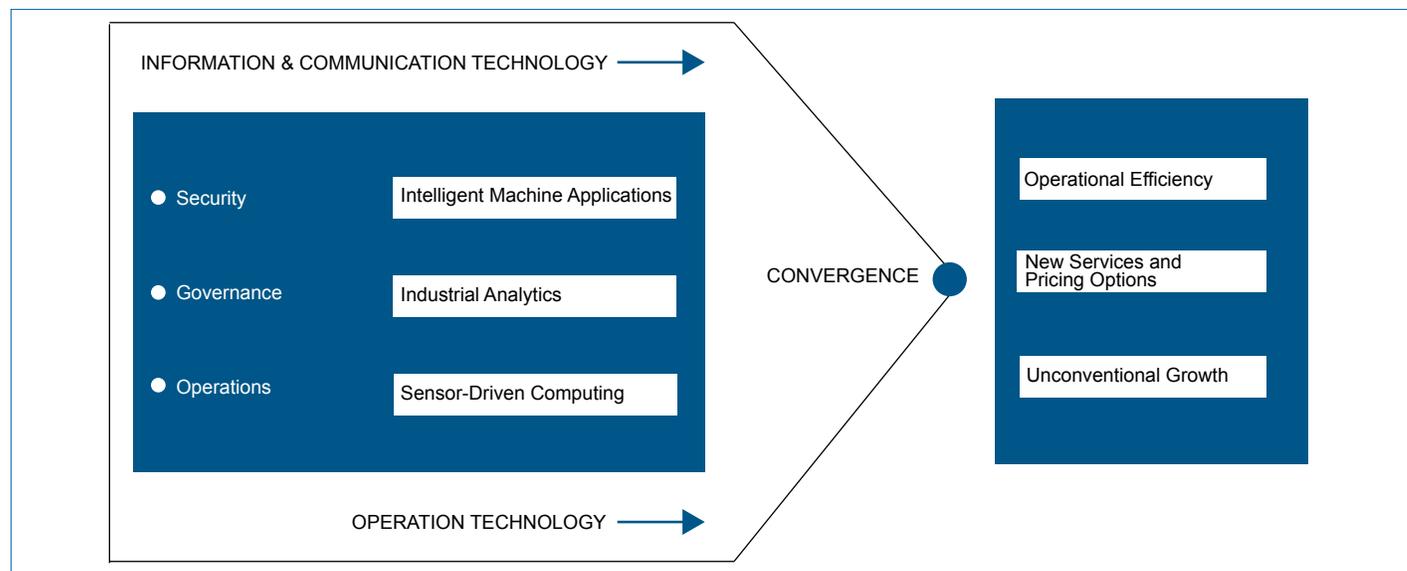
In December 2013, Target Corp’s data breach rendered 40 million customers’ banks accounts compromised. The source of the breach was found to be Fazio Mechanical, a small firm that has commercial relations with Target and whose network had been breached via email malware. The cybercriminals used this network breach to remotely connect to Target’s network. This single and seemingly minuscule attack also managed to affect cash registers in more than 1,800 stores across the United States; it was subsequently found that Target’s computer network was exposed to several vulnerabilities, such as missing patches in the operating system and outdated software, that were easily exploited.

Similarly, in March 2016, investigators at Verizon reported on several breaches against a water utility, referred to using the fake moniker “Kemuri Water Company,” due to what was found to be poor security infrastructure and operational technology systems that were decades old. The SCADA system of the water company, which connected the main operational technology systems (such as valve applications and financial systems), was an IBM AS/400, introduced in 1988. Hackers managed to manipulate the weak system and impede on water treatment and production to the point that the entire process became impaired. Moreover, investigation reports found that the culprits were much less skilled than what one might have expected. According to Verizon’s “Data Breach Digest,” only a small number of security breaches constituted the vast majority of major cyber attacks in a three-year review (Kovacs 2016).

INDUSTRIAL INTERNET OF THINGS

For industry, the Industrial Internet of Things (IIoT) is altering manufacturing, energy, transportation, cities, medical and other industrial sectors, thereby driving a fourth wave of industrial revolution.

The IIoT describes machine-to-machine communications where machines interact and communicate with other machines and objects. These communications result in huge volumes of data that are intelligently generated, processed and analyzed, leading to efficient management. The increasing trend toward the IIoT is transforming

Figure 2: IIoT Convergence of Technology

Source: Author.

industries such as transportation, entertainment, medicine, communications and industrial automation by optimizing operations (Lydon 2014).

In the near future, “the intersection of people, data and intelligent machines will have far-reaching impacts on the productivity, efficiency and operations of industries around the world” (Shekhar 2016). The IIoT presents companies with myriad opportunities to upgrade, offer new services, improve products, increase production, create hybrid business models and enter new markets. To reap the full benefits of the IIoT, organizations will need to excel at exploiting three technology capabilities: sensor-driven computing, industrial analytics and intelligent machine applications.

The IIoT is transforming businesses by:

- optimizing asset utilization;
- reducing operational cost;
- improving worker productivity;
- enhancing worker safety;
- creating new revenue streams;
- improving sustainability; and
- enhancing customer experience (Daugherty et al. 2015).

According to a World Economic Forum (2015) report, examples of the IIoT “include using unmanned aerial vehicles...to inspect oil pipelines, monitoring food safety using sensors, and minimizing workers’ exposure to noise, chemicals and other hazardous gases, especially in traditional heavy industries like oil and gas, manufacturing and chemicals.” In the United Kingdom,

a provider of drinking and waste-water services “is using sensors, analytics and real-time data to anticipate equipment failures and respond more quickly to critical situations, such as leaks or adverse weather events” (ibid.) (see Figure 2).

ICSs

Initially, “ICSs had little resemblance to traditional information systems” as “they were isolated systems running proprietary software and control protocols” (Ross et al. 2006). As these systems have increasingly been integrated “into mainstream organizational information systems to promote connectivity, efficiency and remote-access capabilities, they have started to resemble traditional information systems” and “in many cases, ICSs are using the same commercially available hardware and software components as in the organization’s traditional information systems” (ibid.). According to the report by Ron Ross et al. (2006), “While the change in ICS architecture supports new information system capabilities, it also provides significantly less isolation for these systems from the outside world and introduces many of the same vulnerabilities that exist in current networked information systems. The result is a greater need to secure ICSs.”

There are several drawbacks to traditional ICSs:

- Software, sensors and controls running many contemporary facilities and equipment are outdated and difficult to upgrade. Thus, organizations cannot readily incorporate new features and improvements.
- There is limited integration between internal systems (such as managerial apps, plant data sources) and external partners, which creates data silos.

- Aging operating systems and vulnerable operational technologies pose security risks because they cannot be easily retired or replaced.
- There is limited embedded computing or intelligence control at the device, product or plant level. (Daugherty et al. 2015)

Previous SCADA systems “took advantage of developments and improvement in system miniaturization and Local Area Networking (LAN) technology to distribute processing across multiple systems.” According to Edvard Csanyi (2013), “the distribution of individual SCADA system functions across multiple systems provided more processing power for the system as a whole than would have been available in a single processor.... Distribution of system functionality across network-connected systems served not only to increase processing power, but also to improve the redundancy and reliability of the system as a whole” (ibid.).

Traditional ICSs use an open-system architecture rather than a vendor-controlled, proprietary environment. They use Internet Protocol (IP) for communication and cloud-based services for agility and lower costs. Newer ICSs “have capabilities to monitor inventories, automatically send emails to order more raw materials, contact shippers of ready to ship product, and track product delivery” (Radvanovsky and Brodsky 2014), for example. Continued evolution of control systems with added emphasis on their cyber security is important and necessary for further automation.

The main advantages of new ICSs are:

- increased output or productivity;
- improved quality;
- increased predictability of quality;
- improved consistency of processes or product;
- reduced labour expenditures; and
- improved safety environment for production and operations (Hayden, Assante and Conway 2014).

RISKS AND CHALLENGES FOR THE IIoT

A great deal can go wrong when manufacturing plants, equipment or remote facilities are interconnected and online, including acute disruptions to operations; remote sabotage and loss of life due to impaired infrastructure; and cyber attacks and data theft by criminals, foreign governments and disgruntled employees (Daugherty et al. 2015). Recently, a floating oil rig’s control systems were hacked and the rig shut down after the saboteurs tilted it, “while another rig became so riddled with computer malware that it took weeks to make it seaworthy again” (ibid.).

It is clear that the IIoT must be underpinned by a well-thought-out cyber/physical-security architecture. This goal can be augmented with the following actions:

- Apply non-invasive techniques to patch remote assets, and use industrial control and automation systems that cannot easily be shut down.
- Manage obsolete and legacy operating systems, hosts and devices that have limited or no security built into them.
- Detect and remediate counterfeit or compromised software and hardware.
- Safeguard the integrity of information and systems so that unauthorized access is detected and data that falls into the wrong hands is not corrupted and then reintroduced into critical processes.
- Control and monitor network connections to ensure that only appropriate ones exist between sensitive industrial equipment.
- Build in fail-safe mechanisms to ensure that compromised IT systems that run ICSs cause no physical harm to people and property, or other severe consequences.
- Understand adversaries’ motivations and adapt risk-mitigation strategies to the main danger, such as one-time theft of records, sabotage or ongoing espionage. (ibid.)

Ernie Hayden, Michael Assante and Tim Conway (2014) list the following challenges presented by new ICSs:

- Security vulnerabilities: An automated system may have a limited level of intelligence and therefore be susceptible to injects that could “confuse” or overwhelm processing capabilities.
- Research-and-development cost: The costs of automating a process may exceed the cost saved by the automation itself.
- High initial cost: The automation of a new product or plant typically requires large initial investment, in particular compared with the unit cost of the product.

CYBER THREATS TO CRITICAL INFRASTRUCTURE

Today, threats to critical infrastructure are increasingly through electronic, radio-frequency or computer-based attacks on the information components that control critical infrastructure. Cyber systems form the central infrastructure of critical sectors, nearly all of which use IT to facilitate core business processes. The cyber systems of

critical infrastructure are thus high-value targets for attack, as disrupting them entails extensive economic, political and social effects.

Numerous kinds of threats exist with varying motivations and capabilities, but all breaches exploit certain kinds of cyber systems of critical infrastructure.

As identified in a study by Nadav Morag (2014), computer systems are generally vulnerable to six types of risk:

- risks due to IT (hardware, software, people, processes);
- risks due to interconnection with outside parties and providers (banks, other companies and so on);
- risks due to outside suppliers (cloud providers, subcontractors and so on);
- risks due to disruptions in IT equipment and logistics;
- new technologies (such as the IoT); and
- threats to upstream infrastructure (power supply, water supply and so on).

There has been a dramatic shift toward engaging computer systems with various types of hardware (i.e., the IoT) — for instance, wireless cardiac pacemakers — rendering further vulnerability. Evolving risk areas include the disruption of cloud infrastructure; physical attacks; criminal data mining; digital fraud; and hijacking unmanned aircraft, vehicles and the like (drones, automated cars and so on).

All critical infrastructure systems have vulnerabilities that can be exploited through “threat vectors.” Overall, vulnerabilities in critical infrastructure may be divided into two major subgroups: technical and non-technical (Edwards 2004).

Technical vulnerabilities can be basic vulnerabilities or application-based vulnerabilities. The former refers to the vulnerabilities of common Internet protocols. The core protocols such as IP, TCP (Transmission Control Protocol) and HTTP (Hypertext Transfer Protocol) were created and implemented without factoring in security features since the Internet was initially used to serve academic and governmental environments, wherein the users were trusted entities. Much later, security countermeasures were included in Internet protocols as add-ons with the proliferation of the Internet. Therefore, the Internet is still vulnerable to basic attacks, such as DoS, eavesdropping, hoaxing and packet sniffing. Apart from basic protocols, there are a number of applications, including operating systems, that run on top of basic protocols. These application vulnerabilities are exploited by attackers to gain access privileges to remote systems, steal information and interrupt service. Although a generalization, hacktivists and cyber warriors usually exploit basic protocols first,

then application vulnerabilities, while cybercriminals often target application vulnerabilities.

In spite of state-of-the-art security systems — such as digital signatures, cryptography, biometric security, firewalls, intrusion-prevention systems and access-control systems — security breaches have increased over the years due to non-technical vulnerabilities relating to people and processes, and even closed systems are targeted and affected by viruses and worms such as Stuxnet. Security experts say that Stuxnet ultimately infected the closed network of the Natanz nuclear plant in Iran by means of USB thumb drives. The weakest link in cyber security is the human being: although technical countermeasures are vital for the security of critical infrastructure, they will not be as effective without the conducive and enabling behaviour of people and processes. Cyber spies usually exploit people and process vulnerabilities.

THE CYBER SECURITY CRISIS

With Internet-based networks increasingly touching every aspect of an organization, a single vulnerability in the system can cause a catastrophic chain reaction. Traditional organizational perimeters are eroding, and existing security defences are coming under much pressure. Point solutions, such as “antivirus software, IDS, IPS, patching and encryption...remain a key control for combatting today’s known attacks,” even though hackers have found new ways to circumvent these controls (EY 2013, 1).

Although many of the initial cyber incidents impacting control systems were not directed at ICSs, wide-spreading Internet worms found their way into ICS networks through connections, remote access or by way of portable media. However, there have been examples of internal and external actors specifically targeting ICSs by exploiting vulnerabilities, commanding unauthorized actions or changing set points.

The 2015 “Dell Security Annual Threat Report” (Dell 2015) stated attacks against SCADA systems quadrupled from 2013 to 2014. Specifically, Dell saw worldwide SCADA system attacks increase from 91,676 in January 2012 to 163,228 in January 2013, and to 675,186 in January 2014.

Cyber attacks are increasingly a concern because of their catastrophic physical implications. The mysterious 2008 explosion of the majority BP-owned Baku-Tbilisi-Ceyhan pipeline in Turkey was only recently revealed to have been a digital attack. At the time, Baku-Tbilisi-Ceyhan was thought to be one of the most secure pipelines in the world. Still, unidentified hackers infiltrated the pipeline through a wireless network, tampered with the systems and caused considerable physical damage in an explosion.

One of the main examples, and a game changer for many organizations, was Stuxnet. It was credited as a precision

attack causing physical damage to Iranian nuclear centrifuges by directing them to spin out of control while simultaneously playing recorded system values that indicated normal functioning centrifuges during the attack. This targeted sabotage made clear the potential of cyber attacks.

According to Hayden, Assante and Conway (2014, 20), “One of the most touted ICS cyber incidents involved the unauthorized release of sewage as the result of malicious operation.... Cyber incidents that impact or take command of the control system have raised the specter of consequences that are not shared by IT. In 2007, researchers at the Idaho National Laboratory (INL) demonstrated the ability of using cyber techniques to make unauthorized changes in ICS components which could result in physical damage.” In 2012, a group calling itself “Cutting Sword of Justice” conducted an attack on Saudi Aramco, one of the world’s largest oil companies. In a matter of hours, 35,000 computers were partially wiped out or totally destroyed.

In 2013, major South Korean banks and broadcasters were hacked, which resulted in bank clients being unable to withdraw money from ATMs and broadcasters’ frozen computer networks (Sang-hun 2013). The attack is suspected to have originated in North Korea, with a malware known as “DarkSeoul,” which paralyzed networks. At the end of the same year, DarkSeoul struck again, affecting 48,000 computers in South Korea, disrupting network systems and erasing hard disks, and attempting also to penetrate South Korea’s nuclear operator, which was operating 23 nuclear power plants (Kwon 2015). The latter attack was described as a spear-phishing attack, in which unsuspecting employees of the nuclear operator opened maliciously coded documents in emails.

More recently, in December 2015, Russia-based hackers were alleged to have caused power blackouts across Ukraine in the first full-fledged attack on an electricity distribution network. Around two million people went without electricity for several hours, and experts say such cyber attacks could happen almost anywhere (Vallance 2016). Russian attackers began sending phishing emails to power-utility offices in Ukraine at least six months before the attack. The emails contained Microsoft Word documents, which, once opened, installed malware. Firewalls prevented the attacked computers from gaining control of larger systems, but the malware, known as “BlackEnergy 3,” obtained access to passwords and log-in details, through which the hackers were able to launch another attack. Over time, they were able to remotely log into SCADA systems. By December 23, 2015, the attackers were remotely controlling SCADA computers and cut power at 17 substations, also jamming company communications so that engineers had difficulty gauging the extent of the blackout.

While there is a growing threat of cyber attacks on critical infrastructure, equally important is the rise of physical attacks on energy, transportation and communications. For instance, damage to undersea cables could significantly impede transactions such as the Society for Worldwide Inter-bank Financial Telecommunications, which transmits about 15 million messages a day via submarine cables to more than 8,300 banking organizations, securities institutions and corporate customers in 208 countries (Burnett 2011). In 2008, a broken submarine cable caused by a ship attempting to moor in bad weather off the coast of Egypt led to an Internet blackout that left 75 million people with limited Internet access. Phone and Internet traffic were severely reduced across a huge swath of the region, by as much as 70 percent in India, Egypt and Dubai (Johnson 2008).

The potential of both digital and/or physical attack on critical infrastructure, and the prospective cataclysmic consequences of such, should be a wake-up call for governments, industry and organizations. There is an urgent need for public and private entities to be aware of the risks and, further, be proactive in protecting their valuable information, thereby improving system performance, reliability and safety.

RECOMMENDATIONS

To realize the full potential of the expanding IoT, businesses and governments will need to first overcome a number of hurdles. Security and data privacy are the most important given increased vulnerabilities to attacks, espionage and data breaches driven by increasing connectivity and data sharing.

The following actions are required for an accelerated development of the IoT:

- **Share best practices:** “Operational safety and security practices vary greatly across industry domains. It is important to understand and document existing best practices across industries.... This will help identify gaps and requirements for potential innovation, standards or new cybersecurity products” (World Economic Forum 2015).
- **Policies:** Organizations “need clear legal guidelines over data ownership, transfer and usage” to realize the full potential of the IoT. “Governments need to collaborate with each other and with industry to harmonize compliance requirements in data and liability laws.... This will streamline data flows within a jurisdiction and across national borders” (ibid.).
- **Regulations:** For heavily regulated industries, such as utilities and health care, to truly benefit from the IIoT, policy makers “will need to revisit and possibly relax existing regulations to provide more

flexibility and incentives” to drive innovation. “In the utilities industry, governments can now tap into the new power of transparency enabled by the IIoT to encourage more competition, market efficiency and better customer services” (ibid.).

- **Digital infrastructure:** The success of the IIoT “depends heavily on the presence of robust infrastructures, such as ubiquitous broadband connectivity and digital sensors.” As emerging-market countries “continue large construction efforts, like roads, airports, factories and high-density buildings, they can avoid costly retrofitting faced by developed countries by installing state-of-the-art embedded sensors from the outset. These capabilities provide a foundation for smart cities, enabling more efficient use of natural resources and better public safety and citizen services. Industry can help government leaders to prioritize infrastructure investments that can provide long-term strategic benefits to economic growth, social impact and political success” (ibid.).
- **Role of manufacturers:** For the Internet to have a positive impact, there is a need for Internet service to be accessible, affordable, interoperable, secure and resilient. Today, virtually anyone can manufacture a connectable device. There are no standards for developing and incorporating safety aspects into these devices. Developing testing systems for existing industry and future products would help create a resilient ICS. Once manufacturers have made a connectable product — whether hardware or software — it is not enough to apply security as a veneer atop products that have already been manufactured. During the manufacturing process, security must be a built-in aspect of design for both hardware and software. Over time, these should evolve as standards that guarantee a certain level of default security to the systems. In the same way that quality benchmarks guide users to discriminate with respect to features, these security measures need to be a part of the embedded standards. For example, leading software manufacturers (product and custom) already have aspects such as software development life cycle as a standard input. This needs to become far more widespread — ubiquitous, in fact.
- **Role of Computer Emergency Response Teams (CERTs):** CERTs can play a major role in standardizing processes in the connected world. Standards must be developed toward the manufacture of IIoT devices. Developing standard operating procedures for information sharing between governments and industry is also important. Repositories of vulnerabilities and laws should be created to be better prepared for future counter malicious activities against industrial systems. A global platform is one way to bring together industries by involving key

stakeholders across the value chain. It “can help raise the collective security awareness by sharing threat intelligence. It can also ensure a unified industry voice when communicating with governments or agencies involving security” (ibid.).

- **Raise awareness among policy makers.** Many public policy makers are not well informed about the impact the IIoT might have on citizens, industry and governments. There is an urgent need for them to be better versed in the technology, its societal and policy implications (such as data security, privacy, education and employment) and the impact on government services (ibid.).
- **Cyber security practices:** “Comprehensive, yet targeted, situational awareness is critical to understanding the wider threat landscape and how it relates to the organization. Cyber threat intelligence can bring this knowledge” as “it incorporates both external and internal sources of risk, and covers both the present and future while learning from the past” (EY 2015). Regularly rehearsing incident-response capabilities through “table top exercises [and] enacting complex incident scenarios” tests the organization’s capabilities and provides better crisis management (ibid.). Cyber security “should become a standing boardroom issue — a vitally important item on the agenda. The organization’s leadership should understand and discuss how cyber security enables the business to innovate, open new channels to market and manage risk” (ibid.).

CONCLUSION

The integration of the IIoT with critical infrastructure means new growth opportunities for organizations and governments across the world. Although there are technological challenges and important hurdles to overcome, in particular concerning connectivity and security, the emerging technology will transform interoperability and efficiency in the modern world. According to Paul Daugherty et al. (2015, 17), “To be a viable stakeholder as well as partner in the digitally contestable future — and thus generate new revenues, governments and industries need to make the necessary changes.” Of prime importance is ensuring data privacy, cyber security and accessibility to the global commons in order to drive innovation and growth. Knowing that attacks can never be fully prevented, organizations and governments should advance their cyber-threat-detection capabilities so that response to threat of attack is proactive and appropriate. Learning how to stay ahead of cybercrime will allow organizations to exploit the opportunities offered by the digital world, while minimizing exposure to the risks and costs of dealing with them.

WORKS CITED

- Brandis, George. 2014. "Opening Address of the Critical Infrastructure Resilience Conference Melbourne, Victoria." May 22. www.attorneygeneral.gov.au/Speeches/Pages/2014/Second%20Quarter%202014/6June2014-OpeningAddressOfTheCriticalInfrastructureResilienceConference.aspx.
- Burnett, D. R. 2011. "Cable Vision." *Proceedings Magazine*, August. US Naval Institute.
- Cha, Bonnie. 2015. "A Beginner's Guide to Understanding the Internet of Things." Recode.net, January 15. <http://recode.net/2015/01/15/a-beginners-guide-to-understanding-the-internet-of-things>.
- Clarke, Richard A. and Robert K. Knake. 2010. *Cyber War: The Next Threat to National Security and What to Do About It*. New York, NY: HarperCollins.
- Csanyi, Edvard. 2013. "Three generations of SCADA system architectures." Electrical Engineering Portal, April 22. <http://electrical-engineering-portal.com/three-generations-of-scada-system-architectures>.
- Daugherty, Paul, Prith Banerjee, Walid Negm and Allan E. Alter. 2015. "Driving Unconventional Growth through the Industrial Internet of Things." Accenture. www.accenture.com/in-en/_acnmedia/Accenture/next-gen/reassembling-industry/pdf/Accenture-Driving-Unconventional-Growth-through-IIoT.pdf.
- Dell. 2015. "Dell Security Annual Threat Report." <http://8f7ff0b2bdcb95e1cf03-005bbf273b9ee62b153151d15b71b4f0.r40.cf1.rackcdn.com/articles/2015-dell-security-annual-threat-report-white-paper-15657.pdf>.
- Denning, Dorothy. 2015. "The Rise of Hacktivism." *Georgetown Journal of International Affairs*, September 8. <http://journal.georgetown.edu/the-rise-of-hacktivism>.
- Desjardins, J. 2015. *The Cybersecurity Boom*. <http://www.visualcapitalist.com/the-cybersecurity-boom>.
- Edwards, Matthew. 2004. *Critical Infrastructure Protection*. Amsterdam, NL: IOS Press BV.
- EY. 2013. "Security Operations: Centers against cybercrime." www.ey.com/Publication/vwLUAssets/EY_-_Security_Operations_Centers_against_cybercrime/%24FILE/EY-SOC-Oct-2013.pdf.
- . 2015. "Cybersecurity and the Internet of Things." Insights on governance, risk and compliance." March. [www.ey.com/Publication/vwLUAssets/EY-cybersecurity-and-the-internet-of-things/\\$FILE/EY-cybersecurity-and-the-internet-of-things.pdf](http://www.ey.com/Publication/vwLUAssets/EY-cybersecurity-and-the-internet-of-things/$FILE/EY-cybersecurity-and-the-internet-of-things.pdf).
- Farrell, Greg and Patricia Hurtado. 2015. "JPMorgan's 2014 Hack Tied to Largest Cyber Breach Ever." Bloomberg.com, November 10. www.bloomberg.com/news/articles/2015-11-10/hackers-accused-by-u-s-of-targeting-top-banks-mutual-funds.
- Hayden, Ernie, Michael Assante and Tim Conway. 2014. "An Abbreviated History of Automation & Industrial Controls Systems and Cybersecurity." A SANS Analyst Whitepaper. <https://ics.sans.org/media/An-Abbreviated-History-of-Automation-and-ICS-Cybersecurity.pdf>.
- Johnson, Bobbie. 2008. "How one clumsy ship cut off the web for 75 million people." *The Guardian*, February 1. www.theguardian.com/business/2008/feb/01/internationalpersonalfinancebusiness.internet.
- Kovacs, Eduard. 2016. "Attackers Alter Water Treatment Systems in Utility Hack: Report." SecurityWeek.com, March 22. www.securityweek.com/attackers-alter-water-treatment-systems-utility-hack-report.
- Kwon, K. J. 2015. "Smoking Gun: South Korea Uncovers Northern Rival's Hacking Codes." CNN, April 23. <http://edition.cnn.com/2015/04/22/asia/koreas-cyber-hacking/>.
- Lee, Robert M., Michael J. Assante and Tim Conway. 2014. "German Steel Mill Cyber Attack." ICS CP/PE (Cyber-to-Physical or Process Effects) case study paper. SANS Industrial Control Systems, December 30. http://ics.sans.org:https://ics.sans.org/media/ICS-CPPE-case-Study-2-German-Steelworks_Facility.pdf.
- Lydon, Bill. 2014. "Internet of Things: Industrial automation industry exploring and implementing IoT." InTech, March/April. www.isa.org/standards-and-publications/isa-publications/intech-magazine/2014/mar-apr/cover-story-internet-of-things/.
- Morag, Nadav. 2014. "Cybercrime, Cyberespionage, and Cybersabotage: Understanding Emerging Threats." Colorado Technical University, October. www.coloradotech.edu/~media/CTU/Files/ThoughtLeadership/cybercrime-white-paper.ashx.
- Mossberg, Walt. 2014. "SmartThings Automates Your House Via Sensors, App." Recode.net, January 28. www.recode.net/2014/1/28/11622774/smartthings-automates-your-house-via-sensors-app.
- Moteff, John and Paul Parfomak. 2004. "Critical Infrastructure and Key Assets: Definition and Identification." Congressional Research Service Report for Congress, Library of Congress, Washington, DC. October 1. www.fas.org/sgp/crs/RL32631.pdf.

- Prince Trust of India. 2014. "Cyber criminals hack smart fridge to send out spam." *The Economic Times* (Mumbai), January 20. <http://economictimes.indiatimes.com/industry/cyber-criminals-hack-smart-fridge-to-send-out-spam/articleshow/29110789.cms>.
- Radvanovsky, Robert and Jacob Brodsky, eds. 2014. *Handbook of SCADA/Control Systems Security*. Boca Raton, FL: CRC Press.
- Ross, Ron, Stu Katzke, Arnold Johnson, Marianne Swanson, Gary Stoneburner and George Rogers. 2006. "Recommended Security Controls for Federal Information Systems." National Institute of Standards and Technology. July. <http://csrc.nist.gov/groups/SMA/fisma/ics/documents/appendix-i.pdf>.
- Sang-hun, Choe. 2013. "Computer Networks in South Korea Are Paralyzed in Cyberattacks." *The New York Times*, March 20. www.nytimes.com/2013/03/21/world/asia/south-korea-computer-network-crashes.html.
- Shekhar, Sidharth. 2016. "IoT Adoption in India: How to Take It Forward?" PCQuest, June 1. www.pcquest.com/iot-adoption-in-india-how-to-take-it-forward/.
- Vallance, Chris. 2016. "Ukraine cyber-attacks 'could happen to UK.'" BBC.com, February 29. www.bbc.com/news/technology-35686493.
- World Economic Forum. 2015. "Industrial Internet of Things: Unleashing the Potential of Connected Products and Services." www3.weforum.org/docs/WEFUSA_IndustrialInternet_Report2015.pdf.

ABOUT THE AUTHOR

Tobby Simon is the president and founder of Synergia Foundation, a think tank that works closely with academia, industry and government to establish impactful solutions in the areas of geo-economics and geo-security. He is a commissioner with the Global Commission on Internet Governance and a member of the Trilateral Commission. Tobby has a postgraduate degree in management and is a graduate of the Harvard Business School. He was a research affiliate at the Massachusetts Institute of Technology for five years and has served on the international advisory council of the Belfer Center for Science and International Affairs at the John F. Kennedy School of Government, Harvard University.

The Synergia Foundation is a Bangalore-based interdisciplinary think tank that works with industry, government and academia to establish leading-edge practices in the domains of geopolitics, geo-economics and geo-security. The foundation has multidisciplinary teams that pursue non-partisan research and draws on a global network of resources to offer research analysis and solutions.

**CHAPTER EIGHT:
COMBATting CYBER THREATS: CSIRTS AND FOSTERING
INTERNATIONAL COOPERATION ON CYBER SECURITY**

Samantha Bradshaw

Copyright © 2016 by the Centre for International Governance Innovation
and the Royal Institute of International Affairs

ACRONYMS

APCERT	Asia Pacific Computer Emergency Response Team
CERT/CC	Computer Emergency Response Team Coordination Center
CSIRTs	computer security incident response teams
ENISA	European Union Agency for Network and Information Security
FIRST	Forum of Incident Response and Security Teams
IETF	Internet Engineering Task Force
IGF	Internet Governance Forum
IP	Internet protocol
IR	international relations
ISO	International Organization for Standardization
IT	information technology

INTRODUCTION

In 1988, the first computer worm was unleashed. Robert Morris, a 23-year-old student at Cornell University, created a string of code that spread from computer to computer, causing them to consume memory and shut down. Security experts estimated that the worm took down approximately 10 percent of the network at the time (Madnick, Li and Choucri 2009, 2), and although Morris intended no harm, the worm caused thousands of dollars in damage. A team of programmers at Berkeley and Purdue eventually found solutions and stopped the worm. Morris was convicted under the 1986 Computer Fraud and Abuse Act. He was sentenced to “three years’ probation, 400 hours of community service, and fines of US\$10,000” (Horne 2014, 13).¹

In retrospect, the Internet community realized that the information needed to stop the spread of the Morris worm did not get out as quickly as it could have due to a lack of communication and coordination among the experts working to contain the incident. A US Defense Advanced Research Projects Agency panel suggested that “a lack of communication not only resulted in redundant analysis, but also delayed defensive and corrective measures which could have limited the damage done by the worm” (Ruefle et al. 2014, 19). The panel also concluded that a

formal institution was needed to quickly and effectively coordinate communication among experts during similar security events. Seven days later, it contracted the Software Engineering Institute at Carnegie Mellon University to establish the first computer security incident response team (CSIRT) — the Computer Emergency Response Team Coordination Center (CERT/CC) — to facilitate responses to future cyber security incidents (Ruefle et al. 2014).

The cyber threat landscape has evolved considerably since the first worm. In 2014 and 2015, several events occurred: a high-profile hack against Sony; costly data breaches against companies such as Home Depot, eBay and Target; the discovery of a major “zero-day”² vulnerability called Heartbleed; and the detection of new government-sponsored malware families, such as CosmicDuke, Sandworm and Regin. As innovation continues in areas such as cloud computing, mobile applications and the Internet of Things, significant new security challenges are bound to arise. “Smart” technology provides more opportunities and vectors for attack. As it becomes increasingly integrated into the fabric of our social, economic and political lives, there is ever-greater incentive — and opportunity — for certain actors to try to exploit these systems.

The adversaries in cyberspace have also changed. Today’s cyber threat landscape is composed of a diverse array of aggressors, including large-scale criminal enterprises, curious hackers and state-sponsored groups (Horne 2014). The economics of launching cyber attacks favours the attacker (Center for Strategic and International Studies 2014). Aggressors can easily create malware or acquire it at a low cost. Exploits and vulnerabilities are constantly being discovered, and a black market dedicated to selling these discoveries has emerged. The motivations of these actors vary from political protest to trolling the Internet, stealing personal or financial data, stealing intellectual property and damaging critical infrastructure. Unsurprisingly, governments and armed forces view cyberspace as a new battleground, and many have developed sophisticated scripts designed to ferret out information about their adversaries in the name of national security or public safety.

Cyber security incidents can have severe consequences for businesses, including liability and loss of reputation, customer confidence and productivity (Ahmad, Hadgkiss and Ruighaver 2012). Businesses can also suffer direct financial costs as a result of data theft or physical damage to operating equipment such as servers. But cyber security incidents could affect more than profit margins: as society becomes ever more dependent on the Internet, cyber attacks could have “devastating collateral and cascading effects across a wide range of physical, economic and

¹ Today, Robert Morris teaches at the Massachusetts Institute of Technology.

² The term “zero-day” refers to vulnerabilities that have not yet been made publicly known.

social systems” (Nolan 2015, 3). Incidents can also have devastating psychological effects, as demonstrated by the suicides of individuals associated with the leak of Ashley Madison customer details in 2015 (Baraniuk 2015).

As a result, governments and corporations are increasingly attempting to secure cyberspace, and to secure their systems and citizens from threats that originate there. Cooperation around the prevention of and response to cyber attacks has become an integral component of the cyber security policies of governments from around the world and companies from all sectors of the economy. Currently, private actors play an important role in this partnership, as they own the majority of Internet infrastructure and continually work to secure their networks. Nevertheless, the current institutional landscape for managing cyber security incidents is growing (Choucri, Madnick and Ferwerda 2013). It is made up of thousands of actors: network operators and Internet service providers; businesses and vendors; techies; law enforcement agencies; critical infrastructure operators; governments and military institutions; policy makers; diplomats; and lawyers. Each form a key part of the “regime complex”³ emerging in cyberspace (Nye 2014).

CSIRTs⁴ are also key actors. CSIRTs form an independent network of technical experts that “responds to computer security incidents, coordinates their resolution, notifies its constituents, exchanges information with others, and assists constituents with the mitigation of future incidents” (Best Practice Forum 2014, 3). CSIRTs are often thought of as the “firefighters” (Ahmad, Hadgkiss and Ruighaver 2012, 643) or first-line responders of cyberspace. As the threat landscape has evolved, teams have adapted and expanded by forming an “epistemic community” (Haas 1992) that cooperates to protect and enhance the security and resilience of the Internet.

The changing nature of the current cyber threat landscape has created a need not only for specialized skills in the prevention of and response to cyber attacks, but also for cooperation on a global scale. However, cooperation has been extremely difficult to achieve, especially in regard to information sharing among CSIRTs. Teams generally agree that cooperation could be strengthened through the enhanced and timely exchange of cyber threat information. However, a number of complex legal questions and a lack of trust among the community members have discouraged sharing. This chapter examines the role of CSIRTs in the

emerging cyber regime complex and asks what might be driving the lack of trust and information sharing among the community.

This chapter argues that a number of internal coordination challenges and exogenous contextual problems are influencing the institutional dynamics of CSIRTs. These challenges are giving rise to and exacerbating existing problems regarding information sharing and trust. First, the commercialization of cyber security and the commodification of vulnerabilities such as zero-days have contributed to a competitive, rather than collaborative, approach to cyber security. Second, states are increasingly recognizing the Internet as a new domain in which to exert control. Rather than cooperating with each other and with other actors in the emerging cyber regime complex to strengthen the security of the network, state actors are increasingly hoarding their knowledge of vulnerabilities and other threat-related information that could help CSIRTs prevent and respond to incidents. Third, CSIRTs are increasingly becoming enmeshed in the emergence of a broader cyber regime complex. Teams no longer form a single regime of actors operating in an environment characterized by shared norms, beliefs and procedures. Instead, they must operate in a high-stakes environment shared with other institutions and organizations that have their own distinct and sometimes divergent laws, interests and cultural contexts. Finally, the CSIRT community itself is growing. The importance of the Internet and our dependency on it have increased not only the stakes but also the number of players with interests in protecting and securing the network. Thus, not only are new CSIRTs being socialized into the CSIRT community, where they must coordinate with one another, but the CSIRT community is also being socialized into the broader cyber regime complex, where they must cooperate with a broad range of actors who hold diverging interests. Together, these processes are creating a number of challenges for (international) cooperation.

The first section of this chapter will highlight some key attack trends that characterize the current cyber threat landscape. The second section will provide background information on the global CSIRT network, by describing the current roles and responsibilities a CSIRT assumes and exploring current cooperation, collaboration and information-sharing efforts. The third section will focus on the legal obstacles and trust deficits that limit information sharing. The fourth section will explain how different internal coordination challenges and exogenous effects limit information sharing and trust within the community and among actors operating in the emerging cyber regime complex. The fifth section draws on international relations (IR) literature to discuss how trust can be built within the CSIRT community to remedy some of the information-sharing problems. This chapter concludes with a summary of the findings and makes some recommendations for how

3 On regime complexes, see Raustiala and Victor (2004); Betts (2010); Keohane and Victor (2011); Orsini, Morin and Young (2013); and Drezner (2009).

4 Other names used include, but are not limited to, CERT (a trademarked term referring specifically to the Computer Emergency Response Team of the CERT Coordination Center), CSIRC (computer security incident response capability), CIRT (computer incident response team), IRC (incident response centre) and SERT (security emergency response team).

CSIRTs can be leveraged to improve and coordinate the international response to cyber security incidents.

CYBER THREAT LANDSCAPE

We live in a digital information age in which safeguarding the privacy and security of online data has become an increasingly important concern. Between 2010 and 2014, a number of data breaches took place, increasing the visibility of information security concerns in popular media (see Figure 1). CSIRTs play an active role in protecting the privacy and security of data for their constituents, and in helping to respond to such incidents.

Trends in media coverage are a good indicator of an issue’s salience, but such coverage is prone to hype and can exaggerate the relative occurrence of a problem (Silver 2015). Looking at trends in the frequency of detected web-based attacks provides another angle from which to view the issue. Many (though not all) web-based attacks are aimed at stealing data, thus an analysis of the frequency of such attacks can provide a more well-rounded view of the state of information security. Figure 2 provides a snapshot of the frequency of detected web-based attacks as recorded by Kaspersky Lab.

Some research notes that the apparent rise in cyber attacks can be attributed simply to the growing size of cyberspace and the overall increase in activity, users and points of interaction online (Jardine 2015). Nevertheless, even when normalized around the volume of web traffic and the number of Internet users to account for the growth of cyberspace, the frequency of web-based attacks is still worse now compared to the previous decade and closely mirrors the shape of the media analysis indicators. While the media analysis is not reflective of the drop-off in actual web-based attacks, according to Gartner’s Hype Cycle it could still be on the upward trend of the “technology trigger,” where early media coverage triggers significant

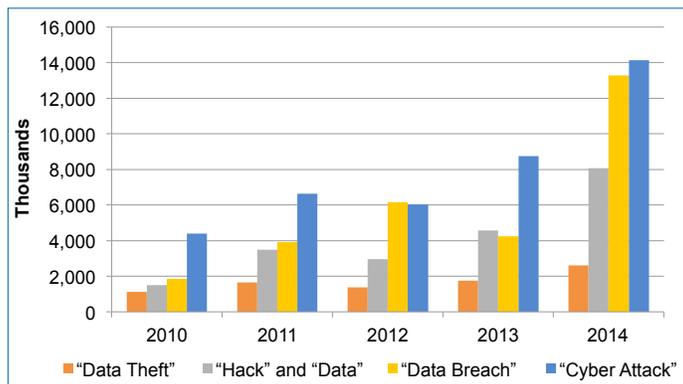
public interest that is not necessarily reflective of the actual occurrence of an event (Gartner 2015). Once people come to recognize the exaggerated nature of the coverage, we can expect such coverage to drop significantly (ibid.; see also Silver 2015).

Nevertheless, people are becoming more cognizant of threats to their own information security. According to a CIGI-Ipsos (2014) poll, which surveyed over 23,326 respondents in 24 countries, 77 percent of users are concerned about someone hacking into their online accounts and stealing their personal information, and 78 percent are concerned about a criminal hacking into their personal bank account.

Yet, despite the fact that people are becoming more aware of their online security and privacy, attackers use “humans more frequently than technology as the weak link” (Ruefle et al. 2014). Hackers and security practitioners refer to this tactic as “social engineering.” Back in the mid-2000s, a phishing prank circulated around the Web where users would receive an email with the subject line “free cup holder.” If the recipient opened the email attachment, a script would open the computer’s CD-ROM drive. While this prank was ultimately harmless, more malicious scripts exploit humans as the weak link in security (Verizon 2015). Today, there has been a surge (or resurgence) of malware that can harvest financial information from victims, record audio or turn on a user’s webcam without their knowledge, record a user’s screen, log keystrokes to steal passwords, or give an attacker remote access to a user’s devices and applications.

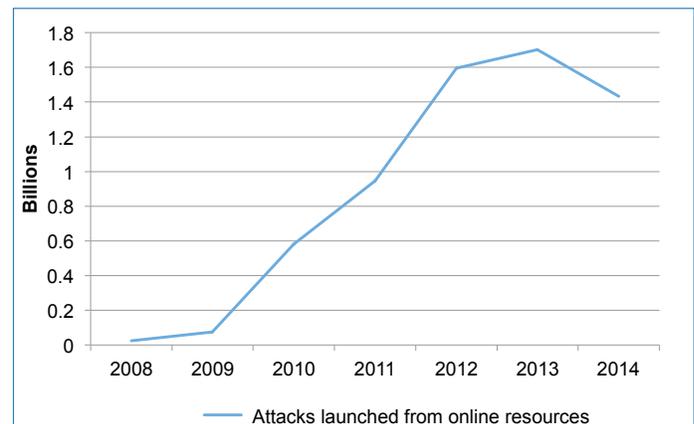
CSIRTs and other cyber security specialists often refer to two broad categories of attacks: targeted and untargeted. Targeted attacks single out an organization or an individual for a specific reason. Targeted attacks take much longer to execute, as an adversary will invest time in finding the best route to deliver an exploit (CERT-UK 2015). One example would be deploying a botnet to deliver a distributed denial

Figure 1: Media Analysis – Information Security Terms 2010–2014



Source: Author; terms listed above were searched in Factiva database from 2010 to 2014.

Figure 2: Frequency of Web-based Attacks



Source: Author; data collected from Kaspersky Lab (2008; 2009; 2010; 2011; 2012; 2013; 2014).

of service attack against a target to overload its network with requests. Another example would be undermining a company's supply chain to corrupt physical equipment or software being delivered to it (ibid.). While they might sound unusual, targeted attacks such as these can be extremely effective and take down some of the most capable organizations. For example, in early 2015 an unprecedented targeted attack against security provider Kaspersky Lab was carried out by attackers who corrupted the digital certificates of software being used by Kaspersky to sign and install a malicious driver on their servers (Zetter 2015). Similarly, in 2008 the US Department of Defense suffered a significant compromise when an infected flash drive was inserted into a US military laptop in the Middle East (Lynn 2010).

In contrast to a targeted attack, untargeted attacks do not discriminate: they will target as many devices, services or users as possible (CERT-UK 2015). Phishing techniques are one type of untargeted attack that involves sending to a large number of people emails that encourage them to give up sensitive information by asking them to reply to an email or open an attachment. Ransomware is another popular method of an untargeted attack. This type of malware prevents users from accessing their system unless they pay the creators a ransom.

Cryptolocker was one ransomware variant that was believed to have been created by a Russian cybercriminal group. It encrypted files on Windows and was believed to infect more than 500,000 victims who were presented with a demand to pay US\$400 within 72 hours or have the keys to their encrypted files destroyed (Ward 2014). In the summer of 2014, CSIRT teams from FireEye and FOX-IT were able to reverse-engineer the Cryptolocker code, and launched a free portal that victims could use to unlock their encrypted information. Despite the success in reducing Cryptolocker, new variants of the malware continue to proliferate on the Web.

It is important to note that the distinction between targeted and untargeted attacks is not always clear and that these techniques can be used in conjunction with one another. Sometimes untargeted attacks are used to carry out targeted ones. An attack by Lizard Squad is one example of this phenomenon. Attackers first compromised thousands of small- and home-office routers with malware. Once they achieved a large enough attack platform, they targeted specific organizations, such as Sony's PlayStation Network and Xbox Live (Passary 2015).

Attackers also take advantage of vulnerabilities in software. An entire market has materialized to sell recently discovered software vulnerabilities that are not yet publicly known — zero-days. Once a zero-day is public, reusable attacks that exploit these vulnerabilities are developed and become openly available (CERT-UK 2015). For example, one study found 85,000 different malware variants that exploited recently publicized zero-days, posing a huge risk to any

device not patched with a security update (Bilge and Dumitras 2012). This problem is further exacerbated by the fact that security patch development and adoption by users can be relatively slow, increasing the window for an attacker to exploit an end user.

The cyber security challenges posed by vulnerabilities are certain to increase for the foreseeable future. With the Internet of Things, there is more potential for vulnerabilities to be discovered and exploited. When everything is a part of the Internet, individuals might not be aware of the fact that their, say, light bulbs and toothbrushes need to be patched and updated. All that is needed from an attacker is an entry point into the network, and the Internet of Things vastly increases the number of vectors for attack as well as the overall size of the attack surface.

In today's cyber threat landscape, a wide variety of skills and coordination are needed to combat increasingly complex challenges. CSIRTs are essential actors with the technical skills necessary to provide incident response and prevention within this changing environment. Given the transnational nature of cyber attacks and the current threat landscape, CSIRTs have formed an informal network to cooperate in preventing and responding to such attacks. The following section details the history, roles and responsibilities of CSIRTs in more detail and discusses current cooperation efforts in the emerging cyber regime complex.

CSIRTs

CSIRTs are teams of experts that use their specialized skills and knowledge to prevent, detect and respond to security incidents for the broader Internet community. Teams form a "global network,"⁵ coming from a diverse group of organizations and institutions, including private sector organizations such as banks and Internet service providers, governments and technical organizations. The roles of various CSIRTs are also diverse, and differ based on factors such as their constituency, skill set and funding levels. This chapter breaks down the classification of teams into three major categories,⁶ based on the parent organization. These categories are:

- **National CSIRTs:** National CSIRTs are the national point of contact for incident response. Broadly speaking, they carry out certain aspects of a state's cyber defence policy — usually by issuing various alerts and warnings, handling aspects of cyber incidents or providing training and education to government constituents. Some national CSIRT

5 For more on global governance networks, see Slaughter (2006); see also Ansell, Sondorp and Stevens (2012).

6 There are many different ways to classify CSIRTs. Some organizations classify them based on the services they provide, their constituency or their parent organization. For an overview of different CSIRT classifications, see Skierka et al. (2015, 12).

capabilities are very advanced and are part of a larger national security operations centre; others are less developed and operate within a particular government department such as law enforcement, military or the ministry of technology or telecommunications. In some countries, more than one national CSIRT exists. Examples of national CSIRTs include the CERT Coordination Centre of Korea, the Canadian Cyber Incident Response Centre, CERT-SE of Sweden and the Chilean Computer Emergency Response Team.

- Private CSIRTs:** These CSIRTs operate for or within a private organization and respond to incidents for their defined constituents. Private CSIRTs could serve a company internally, such as a bank, Internet service provider, or a chemical or petroleum company, or they could be a public-facing for-profit vendor that sells CSIRT services to individuals or companies that do not have in-house security functions. Private CSIRTs can also operate across private companies or across a particular industry category such as banking or e-commerce. Examples of private CSIRTs include the Amazon Security Incident Response Team, the Financial Services Information Sharing and Analysis Centre, the Canadian Imperial Bank of Commerce Incident Response Team, the Symantec CERT and the Verizon CSIRT.
- Technical or Academic CSIRTs:** CSIRTs in this category serve a university or a technical organization, or promote research, education and information sharing within a non-governmental organization. Examples include the Internet Corporation for Assigned Names and Numbers CIRT, the CERT/CC and the Oxford University CERT. Regional organizations such as Asia Pacific CERT (APCERT) or Africa CERT are also included in this category.

Typically, the CSIRT’s constituency will fund the team, determining who it provides services to as well as the kinds of services it will offer. However, some CSIRTs are funded by other organizations or institutions. For example, CGI.br provides CSIRT services to the government of Brazil, but it is not a national CSIRT. To maintain this independence, CGI.br receives its funding from domain name registration in Brazil (Best Practice Forum 2015).

Many view a CSIRT’s role as purely reactive. However, this view does not capture the range of a CSIRT’s capabilities. Isabel Skierka, Robert Morgus, Mirko Hohmann, and Tim Maurer (2015, 13) have noted that “[w]hile the name ‘Computer Security Incident Response Team’ suggests a focus on ‘response,’ CSIRTs provide a range of services.” In addition to reactive services, many teams adopt proactive roles, by, for example, developing security tools, performing risk analysis and testing products for vulnerabilities, providing education to employees on security matters, and operating information security bulletins to share

important information pertaining to vulnerabilities and software patches. However, these kinds of proactive roles tend to only be adopted by more mature CSIRTs (Pereira 2015). Figure 3 provides an overview of various proactive, reactive and security management services a CSIRT can provide to its constituency.

Although teams come from a wide background and have varying levels of skills, the CSIRT community is loosely coordinated through one global organization, the Forum of Incident Response and Security Teams (FIRST). FIRST was founded in the United States in 1990 with the mission of improving information sharing and assisting in the coordination of CSIRTs during network-wide incidents.

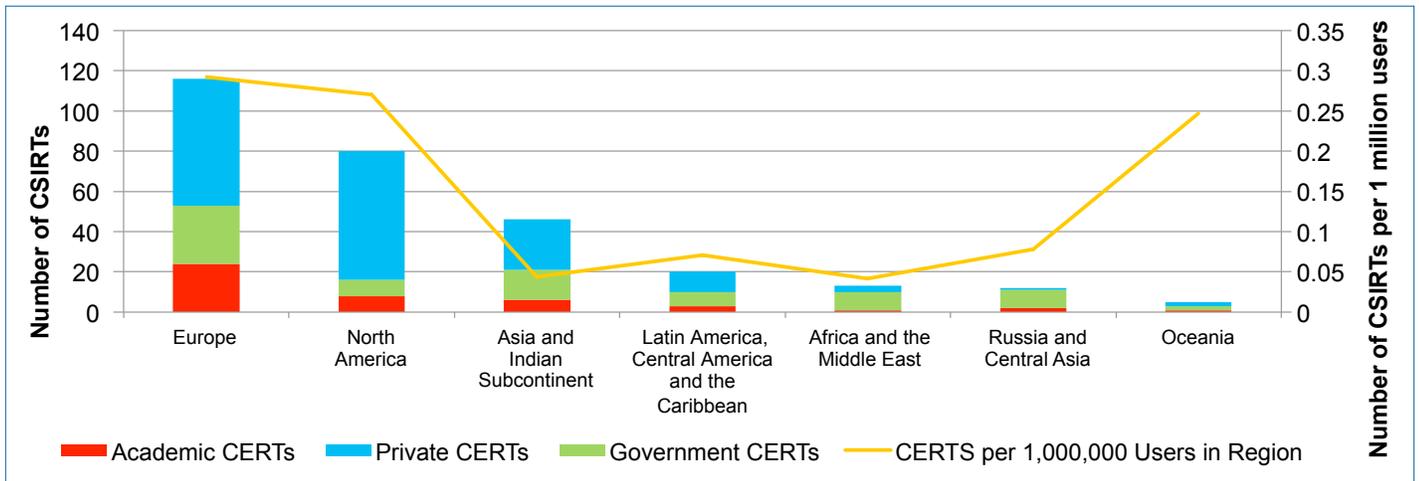
On a global level, FIRST aims to foster cooperation and coordination in incident prevention, to stimulate rapid reaction to incidents and to promote information sharing among members and the community at large. FIRST also plays a large role in promoting best practices and standards for cyber security. It works with other international organizations, such as the International Telecommunication Union and the International Organization for Standardization (ISO), and develops curricula to build and strengthen CSIRT capacity and maturity (FIRST.org 2015).

Currently, there are over 300 CSIRTs from around the world registered with FIRST. The teams come from government, the private sector and academia. They are also geographically diverse, although representation from Africa, the Middle East and Latin America is limited (see

Figure 3: CSIRT Services

Reactive Services	Proactive Services	Security Quality Management Services
Alerts and warnings	Announcements	Risk analysis
Incident handling <ul style="list-style-type: none"> Analysis Response on site Response support Coordination 	Technology watch	Business continuity and disaster recovery planning
Vulnerability handling <ul style="list-style-type: none"> Analysis Response Coordination 	Security audits or assessments	Security consulting
Artifact handling <ul style="list-style-type: none"> Analysis Response Coordination 	Configuration and maintenance of security tools, applications and infrastructures	Awareness building
	Development of security tools	Education/training
	Intrusion detection services	Product evaluation or certification
	Security-related information dissemination	

Source: CERT.org. “Incident Management — CSIRT Services — Service Categories.” www.cert.org/incident-management/services.cfm. Reprinted with permission.

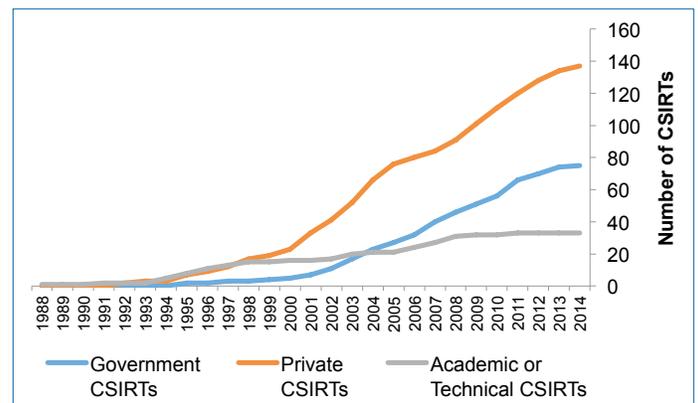
Figure 4: FIRST Membership CSIRT Composition by Region

Source: Bradshaw, Raymond and Shull (2015).

Figure 4). In order to become a FIRST member, CSIRTs need to go through a community validation process.⁷ Once a team becomes part of the FIRST community, it can access incident response information, participate in conferences and technical colloquia and exchange best practices.

In many countries, technical or academic CSIRTs were the first to emerge. As the Internet grew commercially, private companies and government agencies began creating their own teams (see Figure 5). Today, private sector CSIRTs make up the majority of teams and are seen as playing a more direct role in Internet security, due to their role in manufacturing hardware and software and in ensuring their products remain updated and secure. The community views private sector CSIRTs as able to provide “detailed skills and capability in a more narrow topic [compared to] a national CSIRT, which has to respond to incidents across a far more heterogeneous network” (Best Practice Forum 2014, 9).

Due to their direct role in cyber security, private sector CSIRTs also play an important role in international cooperation, knowledge sharing and capacity building by adopting or promoting certain global cyber security standards, sharing information about threats or participating in organizations such as FIRST. However, the Internet’s rapid growth and its importance around the world have highlighted the need for all geographic regions to strengthen their cyber security policies and capabilities through government cooperation. Accordingly, a number of states have worked to develop national CSIRT capabilities. Skierka and colleagues (2015, 8) note that “the expanding role of the state in the governance of CSIRT activities is part of a broader process wherein governments

Figure 5: FIRST Membership Growth 1988–2014

Source: Author; data collected from FIRST.org.

increase regulation of and oversight over the information and communications technology sector.”

Finally, in addition to global organizations such as FIRST, regional and service-specific mechanisms exist that help CSIRTs share knowledge, strengthen capacity and cooperate. These organizations include the European Union Agency for Network and Information Security (ENISA) and Trusted Introducer, which help facilitate knowledge exchange and collaboration among European CSIRTs; APCERT, which coordinates CSIRT organizations in Asia; the Internet Engineering Task Force (IETF); and ISO, which provides standards for CSIRT services and security management.⁸

No matter how strong one’s cyber defence, there is no guarantee that intrusions or incidents will not occur.

⁷ More information on the validation process is detailed on the FIRST website: www.first.org/members/application.

⁸ There are many other regional and service-specific organizations that help facilitate CSIRT cooperation. For more information see ENISA (2013); see also Bada et al. (2014).

CSIRTs play important preventative and responsive roles in cyber security. Although the community is loosely networked, achieving rapid coordination among hundreds of independent entities seems unlikely for a number of reasons. The following section explores some of the information-sharing and trust challenges facing the community.

In addition to reviewing the literature on cyber security cooperation, the following section draws on interviews conducted with CSIRT members who attended the 2015 annual FIRST conference, to provide their detailed insight into perceived cooperation challenges. The forum took place June 14–19, 2015, in Berlin, Germany. It brought together more than 800 leading information technology (IT) experts and practitioners from the security operations community to share knowledge and best practices, to build capacity and to strengthen trust among each other. Conference participants came from around the world, with representation from North America, Latin America, Asia, Africa and Europe.

INFORMATION SHARING AND TRUST DEFICITS

There is general agreement in the CSIRT community that cooperation could be strengthened through the enhanced and timely exchange of cyber threat information between government, private, and academic or technical teams. Information sharing can happen informally in person or by virtual means, or formally through various platforms. Some platforms require members to be from a particular sector or constituency, such as the Information Sharing and Analysis Centre, used to share cyber-related information among incident responders working in the financial sector, or the Cyber Information Sharing and Collaboration Program, used to share information among critical infrastructure operators. There are also a number of free and open-source platforms for information sharing that anyone can refer or contribute to.

The argument for sharing threat-related data is grounded in the belief that cyber security depends on timely and actionable information about threats and the strategies needed to successfully combat them. Information about threats can improve an organization's situational awareness, expand its understanding of the current threat horizon and increase its defensive agility by improving decision making (Ruefle et al. 2014). By leveraging the "capabilities, knowledge and experience of the broader community," organizations can enhance their own cyber defences (Zheng and Lewis 2015).

Threat-related information — such as Internet protocol (IP) or email addresses — is essential for the work of the CSIRT. By identifying and blocking certain addresses known to commit fraudulent phishing schemes, distribute malware,

host illegal content or deliver a distributed denial of service attack, CSIRTs help stop current attacks and prevent future ones against their constituencies. By learning from the experiences of other CSIRTs, teams can identify and stop these threats more quickly, limiting the damage done. Working in collaboration with law enforcement agencies and governments, they can share this information to help dismantle the networks of cybercriminals.

However, it is important to note that information sharing is not a universal remedy for all types of cyber threats. Oftentimes, humans are the weak link in security, and no amount of information sharing can prevent an incident if an individual is used as the vector for attack. In addition, for many new threats, sophisticated actors create and deploy novel techniques. In the first instance of responding to a new threat, some argue, information sharing is not very useful, because analysts have never encountered that threat before (Rosenzweig 2015). Therefore, the lessons learned from community sharing will be largely inapplicable. However, sharing threat data still remains critical for the overall resilience of the network. There is always the chance that a novel attack has similar characteristics to something the community has seen before, and — even if the attack is purely novel — by improving coordination among the collective community, information sharing can reduce the likelihood of a new threat spreading.

Many cyber security analysts believe that threat intelligence can help prevent or minimize the consequences of an attack. In a survey of almost 700 IT and security practitioners, 80 percent of survey participants who experienced a material security breach during the past 25 months believed that "threat intelligence could have prevented or minimized the consequences of the attack" (Ponemon Institute 2015, 2). Yet, despite the widespread perceived benefits of information sharing, there are a number of legal obstacles that dissuade organizations from sharing the necessary information to make cyberspace more resilient. While all of this legislation serves a very important role in society, regulators need to be mindful of the extent to which laws might hinder the ability of the CSIRT community to secure cyberspace, and to carefully consider the intricacies involved in incident response when drafting, interpreting and enforcing laws.

If a cyber security incident is disclosed, corporate legal teams might have to face a variety of liability cases or civil fines. This problem is especially pronounced in the private sector, as one team member stated: "In addition to the potential reputational and financial damage associated with compromise, corporate legal teams often carefully control, manipulate or otherwise impede the release of breach data because of fear of liability."⁹ In another survey of IT security practitioners, over half of the 700 respondents

⁹ Interview conducted by author, June 17, 2015.

listed worries about the “potential liability [from] sharing” as the main reason for not participating in an initiative for exchanging threat information (Ponemon Institute 2014). Liability cases can have a significant economic toll on a company. For example, Target could have faced up to US\$3.6 billion in fines after it revealed that credit card data from its customers was stolen (Williams 2013).

Liability is not the only legal factor dissuading organizations from sharing information. National laws on data exchange and jurisdiction also impact the formal sharing of data with colleague CSIRTs and others working in the security operations community. In recent years, many states have begun enacting “data localization laws” that prevent certain kinds of information from leaving a state’s jurisdiction (Chander and Le 2015). Such limits on information sharing can seriously affect a CSIRT’s ability to respond effectively to incidents. If teams cannot share information outside of their country, they cannot leverage the international community’s knowledge and experience, which are invaluable because cyber threats transcend national boundaries. This restriction can negatively impact a CSIRT’s ability to respond to threats. For example, due to laws that prevent financial information from leaving the legal jurisdiction of Turkey, practitioners noted that CSIRTs in Turkey struggle to effectively and adequately assist their financial sector constituents during cyber security incidents (Internet Governance Forum [IGF] 2014).

Other national laws that relate to freedom of information — where government agencies are required by law to make certain agency records public — can also dissuade teams from sharing threat data. These laws are especially troublesome for teams working in the private sector whose threat intelligence might contain proprietary information. Andrew Nolan (2015) notes that in the United States, sharing threat data that includes proprietary information could waive the sharer’s intellectual property rights under trade secret law. Many countries have trade secret laws that similarly “encourage companies and individuals to invest in collecting information that could help secure competitive advantages in the marketplace” (*ibid.*, 39). In order for trade secret laws to apply, companies must make efforts to maintain the secrecy of information. For example, in the United States, because threat data often contains proprietary information, by voluntarily sharing this data with a third party, companies risk losing any intellectual property rights protection afforded under the US Uniform Trade Secret Act (*ibid.*).

Privacy laws affect when and how it is appropriate for CSIRTs to use and disclose information. CSIRTs will often use data that could constitute personal information to prevent or respond to incidents, such as IP addresses or emails (Cormack 2011). The mitigation of attacks often cannot be accomplished without sharing this kind of information with other CSIRTs or their constituents in order to protect the network and individuals involved in the incident (Best Practice Forum 2014). For example, many

CSIRTs and law enforcement agencies rely on IP addresses to block malicious websites or servers, or use email addresses to track and block spam or phishing attacks. However, privacy is a malleable concept and determining when it is appropriate to use and disclose information to other teams is often unclear and must be done case by case.

Some have suggested that one way to address these privacy concerns would be to sanitize cyber threat data of any proprietary or personal information. However, the process can be time-consuming and requires significant resources, and CSIRT teams have suggested that by the time all identifiers are removed, the information has become obsolete or useless.¹⁰ There is also no guarantee that sanitizing data will protect privacy. Numerous studies have demonstrated that it is very easy to de-anonymize data and identify individuals (for example, see de Montjoye et al. [2015]).

Even in situations where no legal obstacles to sharing information exist, many teams still opt out of sharing threat data with one another. Some members of the CSIRT community attribute this decision to trust deficits.¹¹ In particular, teams might be unwilling to share information about vulnerabilities because it could make their constituents vulnerable to criticism or incur direct financial costs as a result of reputational damage from a security breach disclosure. These fears can severely limit information sharing and cooperation right from the start, as one team member indicated: “No one likes admitting that a breach took place and even without admitting to being compromised asking for help can suggest that something happened.... Others could use this information against you.”¹² Trust that shared information is properly secured and handled delicately is, therefore, a vital element of cooperation.

The fear of reputational damage is not unfounded. If an organization is compromised, publicizing internal vulnerabilities can cause profit losses that outweigh the initial costs of a breach. Target, for example, was reported to have a “62 percent drop in second quarter profits” as a result of the high-profile theft of credit cards in 2014 (Paton 2014). Another American company, USIS, which performs background checks for federal security clearances, suffered severe reputational damage when it suffered a cyber security attack in August 2014, leading to the loss of contracts and more than 2,500 employees (Jayakumar 2014). Because of the high costs that can be associated with a security breach, trust that information will be handled delicately is critically important, especially to private sector constituencies.

What are some of the factors that contribute to and exacerbate problems regarding information sharing

¹⁰ Interview conducted by author, June 18, 2015.

¹¹ Interviews conducted by author, June 15 and 17, 2015.

¹² Interview conducted by author, June 15, 2015.

and trust? The following section describes four such obstacles: the commercialization of cyberspace and the commodification of vulnerabilities; geopolitical power and cyberspace as a new threat domain; the growth of the CSIRT community; and the emergence of a cyber regime complex.

OBSTACLES TO BUILDING TRUST AND SHARING INFORMATION

Cyberspace has often been characterized as a “competitive environment prone to conflict rather than cooperation” (Ito 2014, 2). The emergence of contention in systems of Internet governance has made cooperation extremely difficult (Bradshaw et al. 2015). An array of public and private actors from around the globe are involved in Internet governance (Raymond and DeNardis 2015), and the diversity of actors involved in Internet governance and cyber security with differing interests, values and views of legitimate procedures for how governance should be conducted has increased the potential for deadlocked negotiations (Bradshaw et al. 2015; Raymond and Smith 2014). All of this is moving the cyber regime further away from the original conception of “cyberspace as a shared global resource” that promotes an open and collaborative environment (Ito 2014, 2). Given the transnational nature of cyber risk, having national governments and private organizations both involved in cyber security increases the importance of cooperation. However, a number of internal coordination challenges and exogenous contextual problems influence the institutional dynamics of CSIRTs. These challenges are giving rise to new problems regarding sharing and trust, and intensifying existing ones.

Commercialization of Cyberspace

The commercialization of cyber security and the commodification of vulnerabilities such as zero-days are factors that have contributed to a competitive, rather than collaborative, approach to cyber security. Information sharing within and across organizations has never been perfect; however, the commercialization of cyberspace has exacerbated many information-sharing deficits.

Cyber vulnerabilities have become increasingly valuable commodities, not only for criminals who wish to deliver exploits but for private CSIRTs whose business models are designed to profit by stopping them. Commercial or vendor CSIRTs that sell services might not always want to share information about threats. Threat data and cyber security defence strategies are tremendously valuable to vendor CSIRTs and sharing this kind of information could hurt their bottom line. At the FIRST conference, it was noted that “if you know what the winning lottery numbers are going to be, you aren’t going to share them” (Railton 2015). Usually, competition is a sign of a healthy marketplace, as it leads to better and more differentiated

products and services. However, because there is imperfect information — where vulnerability data is not equally accessible to those trying to stop threats — competition is leading to more insecurity and less trust among those trying to secure the network.

At the same time, as more businesses move online, the commercialization of cyberspace has increased the cost of a breach. More information and data are now uploaded, shared and stored online. More services are offered online and much of an individual’s social and economic life is integrated into the Internet. As a result, companies that operate online have a great deal at stake. If customers lose confidence in the businesses operating online, profits can drop due to reputational damage and liability. Thus, incident responders are under increasing pressure to quickly and quietly respond to threats — an obstacle to information sharing.

New Threat Domain

A second obstacle is the increasing recognition among states that the Internet is a new domain in which to exert control. Rather than cooperating to strengthen the security of the network, state actors are increasingly hoarding information about vulnerabilities and threats that could help CSIRTs prevent and respond to incidents. One practitioner at FIRST noted that “it is not just the bureaucracy or legal obstacles that limit information sharing between CSIRTs and state actors. State actors are increasingly collecting threat information to develop their own malware and deliver exploits for various national security or surveillance purposes. They don’t want to share this information with us because we could stop their exploits.”¹³

State-sponsored malware is not a new phenomenon, as much evidence exists of state actors using various aspects of the Internet and Internet technology to achieve various political or economic goals (DeNardis 2012; DeNardis 2014; Bradshaw and DeNardis 2015). The earliest reported case of government malware dates back to 2001, when FBI agents snuck into a home and installed a script that recorded keystrokes (Mayer 2015). Although the vast majority of malware is criminal, governments also use it to collect intelligence and carry out covert actions against other states (Electronic Frontier Foundation 2015). Thus, sharing intelligence about vulnerabilities could weaken state efforts to exploit them for national security or other purposes.

Growth of the CSIRT Community

A third problem in establishing trust and information sharing is the growth of the CSIRT community itself. The importance of the Internet and our dependency on it has

¹³ Interview conducted by author, June 18, 2015.

increased not only the stakes of the players with interests in protecting and securing the network, but their number. At one time, there was a single CSIRT responding to incidents. Today, there is a cornucopia of teams operating across governments and all sectors of the economy. As the community continues to grow, competition between teams has become a barrier to their cooperation.

A number of governments have begun to establish national CSIRTs to strengthen their own capacity to prevent and respond to cyber threats. Sometimes, governments appoint more than one national CSIRT. In these instances, private or technical CSIRTs might have provided services for a period of time (Best Practice Forum 2014). This trend has led to increased competition and counterproductive results in the form of non-cooperation, as CSIRTs compete to legitimately represent a national constituency.

Emergence of the Cyber Regime Complex

The fourth obstacle is the enmeshing of CSIRTs within a broader, emerging cyber regime complex. Teams no longer form a single regime of actors operating in an environment characterized by generally held norms, beliefs and procedures. The constituencies of various CSIRTs operating in the emerging cyber regime complex have diverging interests, making cooperation extremely difficult. States view the Internet as a new domain, which has led them to develop their own malware and scripts for exploiting other states, and to hoard zero-day vulnerabilities. The quest for geopolitical power and a strategic military advantage over another state's cyber defences is sometimes at odds with the state's responsibility to ensure public safety and secure cyberspace, because developing new exploits or leaving old vulnerabilities unaddressed creates risk in the system.

Similarly, diverging interests arise due to the commercialization of cyber security and the commodification of vulnerabilities. Market competition is increasingly at odds with ensuring cyber security. Sharing threat-related information is necessary for securing cyberspace, but it can also put a constituency at risk because it often involves revealing information about its own insecurities. Thus, the functional interest of CSIRTs — preventing and responding to incidents — is placed at odds with their material interest in protecting their constituencies' assets and reputations.

Finding a solution to these conflicting interests will likely prove difficult in the foreseeable future. As Joseph S. Nye Jr. (2014, 14) notes: "Predicting the future of the normative structures that will govern [the cyber regime complex] is difficult because of the newness and volatility of the technology, the rapid changes in economic and political interests and the social and generational cognitive evolution that is affecting how state and non-state actors understand and define their interests."

States are important contributors to the norms that define regime complexes (Morin and Orsini 2013). However, non-state actors can also perceive and manage problematic relationships among the different actors within a regime complex (Orsini, Morin and Young 2013). In the area of cyber security, CSIRTs could be leveraged as "norm entrepreneurs" that could link the regimes and their competing interests, and "focus efforts on addressing the problem" to make cooperation more likely (Struett, Nance and Armstrong 2013, 94). After all, Peter M. Haas notes (as cited in Cross [2013, 149]) that epistemic communities are "responsible for developing and circulating casual ideas and some associated normative beliefs...thus helping to create...interests and preferences." CSIRTs have already begun this process, by attempting to develop norms for strengthening trust between each other as well as among their constituents. The following section discusses trust-building initiatives and opportunities to strengthen cooperation among CSIRTs.

NORMS FOR STRENGTHENING TRUST

Ensuring cyber security is a shared mission of governments, private companies and the technical community. In order to overcome some of the challenges in information sharing, CSIRTs have attempted to establish nodes of trust across the community. However, trust-building is only one strategy and can mitigate only some of the information-sharing challenges. For example, greater levels of trust will not solve liability or trade secrecy issues. Laws that address these other issues and encourage information sharing have to be developed in tandem with CSIRT efforts to encourage norms around trust.

Nevertheless, trust is important for strengthening relationships between CSIRTs and other actors who are responsible for securing cyberspace. Teams have to trust that sensitive information about breaches and vulnerabilities will be handled with care, and will not be used with ill intent for unrelated or alternative purposes. One well-known model for building trust within the community is sponsorship, where a trusted team advocates on behalf of a new team that wishes to join the community. Personal relationships play an important role within the CSIRT community because of the high standards placed on the technical expertise and the integrity of a team (Skierka et al. 2015). Generally, the sponsorship model works well in small communities, especially when teams are working within the same sector or on similar issues with similar organizational cultures. Some smaller communities have been extremely effective at establishing cooperative environments with liberal information-sharing policies. However, these trust-building models do not work as well for large groups because entry is extremely difficult and, as groups grow, the level of trust and collaboration often diminishes (Ruefle et al. 2014).

CSIRTs frequently describe trust as a “Catch-22” problem, where one needs to have trust in order to gain it.¹⁴ One of the biggest challenges for building initial trust is uncertainty. Teams can be reluctant to share or disclose relevant information that could make them or their constituents more vulnerable or give another CSIRT company an edge in the marketplace. Furthermore, the disclosures of former US National Security Agency contractor Edward Snowden have brought to light the pervasiveness of surveillance activities by state actors, heightening uncertainty over CSIRT involvement in surveillance operations and discouraging cooperation with teams and organizations involved in national cyber security and law enforcement efforts (Best Practice Forum 2015).

Uncertainty about another’s action is viewed as an obstacle to cooperation (Koremenos, Lipson and Snidal 2011, 765). Finding strategies to reduce this uncertainty is key to improving levels of trust. Strategies such as third-party accreditation have been applied to help build trust within larger groups and to remove uncertainty about a team’s capacity, procedures and policies. For example, third-party accreditation organizations, such as Trusted Introducer, list well-known teams and accredit them according to demonstrated and verified levels of capacity and maturity (Trusted Introducer 2015). Other mechanisms, such as the IETF’s “Best Current Practice 21: Request for Comments 2350” (Brownlee and Guttman 1998), recommend that CSIRTs publish information pertaining to their policies and procedures, services offered and scope of operations. If adopted, these requests for comment can act as another mechanism for reducing uncertainty and building trust by increasing the transparency of a CSIRT’s operations.

Accreditation models have been viewed as beneficial for communities with many participants because they not only verify a certain degree of skill but also allow for the creation of smaller subgroups with higher trust levels (ENISA 2015). However, accreditation mechanisms are entirely voluntary — no official international standards or requirements exist. Instead, those teams that choose to apply for accreditation need only fulfill the specific requirements of the individual certifying organization.¹⁵ Furthermore, these mechanisms do not strictly define the intricacies of handling sensitive information. While it would be onerous to define a strict set of requirements that would be appropriate for all incident responders, improving these standards and making them transparent and obligatory would help to reduce uncertainty around incident response. For example, privacy and other data-handling policies that include provisions on data retention, collection and storage could be updated and made a

necessary requirement for teams seeking membership at FIRST.

Another way CSIRTs try to bridge the gap between competing teams is through membership in organizations such as FIRST. Cooperation can occur on the basis of desired membership in a community with a particular set of values and practices (Johnston 2001). Given its role as a global institution for strengthening CSIRT cooperation, FIRST acts as a normatively desirable community with shared values and best practices, as well as with a certain degree of trust among its members.

Although obtaining membership in a particular group might be a necessary condition for creating trust, membership alone is not sufficient. Teams who join FIRST are quickly isolated if they do not contribute to the shared body of knowledge (Grance et al. 2015). Thus, “reciprocity” is also a key element, especially when a new team is joining the community (Skiera et al. 2015, 21).

Cooperation can also emerge in tit-for-tat behaviour (Axelrod 2006). However, tit-for-tat reciprocity should not be seen as “quid pro quo.” As a concept, reciprocity can have two quite distinct meanings. Robert O. Keohane (1986, 4) distinguishes between *specific* reciprocity, where “specified partners exchange items of equivalent value in a strictly delimited sequence” and *diffuse* reciprocity, which is generally viewed as “an ongoing series of sequential actions [that] may continue indefinitely, never balancing but continuing to entail mutual concession within the context of shared commitments and specific values.” Often when teams share information there is an expectation that information will be shared quid pro quo (Railton 2015). However, because sharing cyber threat information is largely dependent on the timing and current experiences of a team, adopting a diffuse definition of reciprocity could help strengthen trust and build more cooperative relationships.

CONCLUSION

The cyber threat landscape has dramatically changed over the past 25 years. Cyber is now largely an “offense-dominated domain” (Nye 2010), skewed in favour of the attacker, wherein adversaries are able to quickly and cheaply find vulnerabilities and develop new techniques for infiltration. But this chapter suggests that it is not only the threat landscape that is changing: new actors are increasingly becoming involved in cyber governance, and CSIRTs are increasingly becoming enmeshed in an emerging cyber regime complex. Not only do teams have to cooperate with their own growing community, but they must also consider the preferences of other institutions and organizations in their work: market preferences are often placed at odds with ensuring cyber security or protecting human rights; similarly, law enforcement or surveillance activities can be placed at odds with

¹⁴ Interviews conducted by author, June 15 and 18, 2015.

¹⁵ For example, Trusted Introducer’s requirements for CSIRT accreditation are laid out online: www.trusted-introducer.org/processes/accreditation.html.

privacy or ensuring cyber security. Further, as CSIRTs become increasingly commercialized or move into new government or bureaucratic domains, it is important that they do not lose the quality of being a “team” (Best Practice Forum 2015). Informal sharing facilitated by normative communities such as FIRST is important for strengthening trust and building ongoing relationships. Amid bureaucratization and commercialization, these kinds of informal relationships could get lost to process and competition.

Bridging the trust deficits that exist within the community is important to enhancing international cooperation on cyber security. Reducing uncertainty by better defining roles and practices, and by redefining expectations when it comes to information sharing, can help to strengthen cooperation between CSIRTs. By being more transparent with their practices surrounding data, CSIRTs can remain a more neutral actor cooperating across constituencies to promote the ongoing stability and security of cyberspace.

As the nature of cyber threats continues to change, CSIRTs with a variety of skills in incident response will be needed to effectively identify and respond to threats. While the number of CSIRTs in the world is growing, these teams vary widely in their stages of development. Cyber incident response capabilities are in their infancy. As more countries and companies recognize the importance of cyber security and incident response, it will become increasingly difficult to find the right candidates. Even now, many practitioners note that attracting good, effective and efficient talent is hard.¹⁶ Along with bridging the increasingly complex trust deficits within the community and the broader cyber regime complex, capacity building and skills training are needed to help CSIRTs remain effective and able to meet new cyber security challenges.¹⁷

The upside of CSIRT capability becoming enmeshed in the broader regime complex is that many of the other elementary regimes have significant material resources, which provides the CSIRT community with an opportunity to strengthen its own capacity. But to leverage this opportunity, CSIRTs will need more than the technical expertise that traditionally accompanies the job. Specifically, teams will need to expand their skills and expertise into new areas such as law, policy and government, and international relations to operate effectively in the emerging cyber regime complex.

¹⁶ Interviews conducted by author, June 18 and 19, 2015.

¹⁷ For more information on CSIRT capacity building and best practices for CSIRT maturity, see ENISA (2013).

WORKS CITED

- Ahmad, Atif, Justin Hadgkiss and A. B. Ruighaver. 2012. “Incident Response Teams — Challenges in Supporting the Organizational Security Function.” *Computers & Security* 31 (5): 643–52.
- Ansell, Chris, Egbert Sondorp and Robert Hartley Stevens. 2012. “The Promise and Challenge of Global Network Governance: The Global Outbreak Alert and Response Network.” *Global Governance* 18: 317–37.
- Axelrod, Robert. 2006. *The Evolution of Cooperation*. Cambridge, MA: Basic Books.
- Bada, Maria, Sadie Creese, Michael Goldsmith, Chris Mitchell and Elisabeth Phillips. 2014. “Computer Security Incident Response Teams (CSIRTs): An Overview.” Oxford: Global Cyber Security Capacity Centre. www.sbs.ox.ac.uk/cybersecurity-capacity/system/files/CSIRTs.pdf.
- Baraniuk, Chris. 2015. “Ashley Madison: ‘Suicides’ Over Website Hack.” BBC News, August 24. www.bbc.com/news/technology-34044506.
- Best Practice Forum. 2014. “Best Practice Forum on Establishing and Supporting Computer Security Incident Response Teams (CSIRT) for Internet Security.” IGF. www.intgovforum.org/cms/documents/best-practice-forums/establishing-and-supporting-computer-emergency-response-teams-certs-for-internet-security/409-bpf-2014-outcome-document-computer-security-incident-response-teams/file.
- . 2015. “Best Practice Forum on Establishing and Supporting Computer Security Incident Response Teams (CSIRT) for Internet Security.” IGF. www.intgovforum.org/cms/187-igf-2015/transcripts-igf-2015/2324-2015-11-11-bpf-establishing-and-supporting-computer-security-incident-response-teams-csirts-workshop-room-6.
- Betts, Alexander. 2010. “The Refugee Regime Complex.” *Refugee Survey Quarterly* 29 (1): 12–37.
- Bilge, Leyla and Tudor Dumitras. 2012. “Before We Knew It: An Empirical Study of Zero-Day Attacks in the Real World.” Presentation at the 19th ACM Conference on Computer and Communications Security, Raleigh, NC, October 16–18. https://users.ece.cmu.edu/~tdumitra/public_documents/bilge12_zero_day.pdf.
- Bradshaw, Samantha and Laura DeNardis. 2015. “The Politicization of the Domain Name System: Implications for Internet Security, Stability and Freedom.” Paper presented at the European Consortium of Political Research, Montreal, QC, August 29.

- Bradshaw, Samantha, Laura DeNardis, Fen Hampson, Eric Jardine and Mark Raymond. 2015. *The Emergence of Contention in Global Internet Governance*. Global Commission on Internet Governance Paper Series No. 17. Waterloo, ON: CIGI. www.cigionline.org/publications/emergence-of-contention-global-internet-governance.
- Bradshaw, Samantha, Mark Raymond and Aaron Shull. 2015. "Rule Making for State Conduct in the Attribution of Cyber Attacks." In *Mutual Security in the Asia-Pacific: Rules for Australia, Canada and South Korea*, edited by Kang Choi, James Manicom and Simon Palamar, 153–71. Waterloo, ON: CIGI; Seoul, Korea: Asan Institute for Policy Studies.
- Brownlee, N. and E. Guttman. 1998. "Expectations for Computer Security Incident Response." Best Current Practice 21: Request for Comments 2350. IETF, June. www.ietf.org/rfc/rfc2350.txt.
- Center for Strategic and International Studies. 2014. "Net Losses: Estimating the Global Cost of Cybercrime." Center for Strategic and International Studies, June. www.mcafee.com/ca/resources/reports/rp-economic-impact-cybercrime2.pdf.
- CERT-UK. 2015. "Common Cyber Attacks: Reducing the Impact." www.gov.uk/government/uploads/system/uploads/attachment_data/file/400106/Common_Cyber_Attacks-Reducing_The_Impact.pdf.
- Chander, Anupam and Uyen Le. 2015. "Data Nationalism." *Emory Law Journal* 64 (3): 677–739.
- Choucri, Nazli, Stuart Madnick and Jeremy Ferwerda. 2013. "Institutions for Cyber Security: International Responses and Global Imperatives." *Information Technology for Development* 20 (2): 96–121.
- CIGI-Ipsos. 2014. Global Survey on Internet and Trust. www.cigionline.org/internet-survey.
- Cormack, Andrew. 2011. "Incident Response and Data Protection." Version 2. www.terena.org/activities/tf-csirt/publications/data-protection-v2.pdf.
- Cross, Mai'a K. Davis. 2013. "Rethinking Epistemic Communities Twenty Years Later." *Review of International Studies* 39: 137–60.
- de Montjoye, Yves-Alexandre, Laura Radaelli, Vivek Kumar Singh and Alex "Sandy" Pentland. 2015. "Unique in the Shopping Mall: On the Reidentifiability of Credit Card Metadata." *Science* 347 (6221): 536–39.
- DeNardis, Laura. 2012. "Hidden Levers of Internet Control." *Information, Communication & Society* 15 (5): 720–38.
- . 2014. *The Global War for Internet Governance*. New Haven, CT: Yale University Press.
- Drezner, Daniel W. 2009. "The Power and Peril of International Regime Complexity." *Perspectives on Politics* 7 (1): 65–70.
- Electronic Frontier Foundation. 2015. "State-Sponsored Malware." www.eff.org/issues/state-sponsored-malware.
- ENISA. 2013. "CERT Community — Recognition Mechanisms and Schemes." www.enisa.europa.eu/activities/cert/support/baseline-capabilities/cert-community-recognition-mechanisms-and-schemes.
- . 2015. "Models of Trust." www.enisa.europa.eu/activities/cert/background/coop/models-legal/trust-models.
- FIRST.org. 2015. "Standardization Efforts." www.first.org/global/standardisation.
- Gartner. 2015. "Gartner Hype Cycle." www.gartner.com/technology/research/methodologies/hype-cycle.jsp.
- Grance, Timothy, Thomas Millar, Pawel Pawlinski, Luc Dandurand and Sarah Brown. 2015. "Threat Information Sharing: Perspectives, Strategies and Scenarios." Presentation at 27th Annual FIRST Conference, Berlin, June 15.
- Haas, Peter M. 1992. "Epistemic Communities and International Policy Coordination." *International Organization* 46 (1): 1–35.
- Horne, Bill. 2014. "On Computer Security Incident Response Teams." *IEEE Security & Privacy* (September/October).
- IGF. 2014. "BPF3 — Establishing and Supporting CERTs for Internet Security." YouTube video, 1:43:06. Streamed live on September 4. <https://m.youtube.com/watch?v=YnOljPgfqmI>.
- Ito, Yuri. 2014. "The Cyber Green Initiative: Improving Health Through Measurement and Mitigation." JP CERT Coordination Centre, November 17. www.jpccert.or.jp/research/GreenConcept-20141117_en.pdf.
- Jardine, Eric. 2015. *Global Cyberspace Is Safer Than You Think: Real Trends in Cybercrime*. Global Commission on Internet Governance Paper Series No. 16. Waterloo, ON: CIGI. www.cigionline.org/publications/global-cyberspace-safer-you-think-real-trends-cybercrime.
- Jayakumar, Amrita. 2014. "USIS Cuts More Than 2500 jobs after Losing Contracts in Wake of Cyberattack." *The Washington Post*, October 7. www.washingtonpost.com/business/capitalbusiness/usis-cuts-more-than-2500-jobs-after-losing-contracts-in-wake-of-cyberattack/2014/10/07/5816cfb2-4e3f-11e4-babe-e91da079cb8a_story.html.

- Johnston, Alastair Ian. 2001. "Treating International Institutions as Social Environments." *International Studies Quarterly* 45: 487–515.
- Kaspersky Lab. 2008. "Kaspersky Security Bulletin 2008." <http://securelist.com/analysis/kasperskysecuritybulletin/36241/kaspersky-security-bulletinstatistics-2008>.
- . 2009. "Kaspersky Security Bulletin 2009." <http://securelist.com/analysis/kaspersky-securitybulletin/36284/kaspersky-security-bulletin-2009-statistics-2009>.
- . 2010. "Kaspersky Security Bulletin 2010." <http://securelist.com/analysis/kaspersky-securitybulletin/36345/kaspersky-security-bulletin-2010-statistics-2010>.
- . 2011. "Kaspersky Security Bulletin 2011." <http://securelist.com/analysis/kaspersky-securitybulletin/36344/kaspersky-security-bulletinstatistics-2011/>.
- . 2012. "Kaspersky Security Bulletin 2012." <http://securelist.com/analysis/kaspersky-securitybulletin/36703/kaspersky-security-bulletin-2012-theoverall-statistics-for-2012>.
- . 2013. "Kaspersky Security Bulletin 2013." http://media.kaspersky.com/pdf/KSB_2013_EN.pdf.
- . 2014. "Kaspersky Security Bulletin 2014." <http://cdn.securelist.com/files/2014/12/Kaspersky-Security-Bulletin-2014-EN.pdf>.
- Keohane, Robert O. 1986. "Reciprocity in International Relations." *International Organization* 40 (1): 1–27.
- Keohane, Robert O. and David G. Victor. 2011. "The Regime Complex for Climate Change." *Perspectives on Politics* 9 (1): 7–23.
- Koremenos, Barbara, Charles Lipson and Duncan Snidal. 2001. "The Rational Design of International Institutions." *International Organization* 55: 761–99.
- Lynn, William J., III. 2010. "Defending a New Domain: The Pentagon's Cyberstrategy." *Foreign Affairs* (September/October). www.foreignaffairs.com/articles/united-states/2010-09-01/defending-new-domain.
- Madnick, S., X. Li and N. Choucri. 2009. "Experiences and Challenges with Using CERT Data to Analyze." Massachusetts Institute of Technology Engineering Systems Division Working Paper Series. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1478206.
- Mayer, Jonathan. 2015. "Constitutional Malware." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2633247&download=yes.
- Morin, Jean-Frederic and Amandine Orsini. 2013. "Regime Complexity and Policy Coherency: Introducing a Co-adjustments Model." *Global Governance* 19 (1): 41–53.
- Nolan, Andrew. 2015. *Cybersecurity and Information Sharing: Legal Challenges and Solutions*. Congressional Research Service Report. CRS, March 16. www.fas.org/sgp/crs/intel/R43941.pdf.
- Nye, Joseph S., Jr. 2010. "Cyber Power." Belfer Center for Science and International Affairs, Harvard Kennedy School, May. <http://belfercenter.ksg.harvard.edu/files/cyber-power.pdf>.
- . 2014. *The Regime Complex for Managing Global Cyber Activities*. Global Commission on Internet Governance Paper Series No. 1. Waterloo, ON: CIGI. www.cigionline.org/publications/regime-complex-managing-global-cyber-activities.
- Orsini, Amandine, Jean-Frederic Morin and Oran Young. 2013. "Regime Complexes: A Buzz, a Boom or a Boost for Global Governance?" *Global Governance* 19 (1): 27–39.
- Passary, Anu. 2015. "PSN and Xbox Live Go Down: Lizard Squad to Blame?" *Tech Times*, May 16. www.techtimes.com/articles/53512/20150516/psn-and-xbox-live-go-down-and-lizard-squad-takes-credit.htm.
- Paton, Elizabeth. 2014. "Cyber Attack Takes Toll on Target." *Financial Times*, August 20. www.ft.com/cms/s/0/1fcf4c82-287f-11e4-8bda-00144feabdc0.html#axzz3eTdPPUX8.
- Pereira, Nishan Marc. 2015. "The Incident Prevention Team: A Proactive Approach to Information Security." Master's thesis, Delft University of Technology. <http://repository.tudelft.nl/view/ir/uuid%3A21c6b579-a25b-4395-ba88-786e5f1eb33c/>.
- Ponemon Institute. 2014. *Exchanging Cyber Threat Intelligence: There Has to Be a Better Way*. Ponemon Institute Research Report. April. <http://content.internetidentity.com/acton/attachment/8504/f-001b/1/-/-/-/-/ Ponemon%20Study.pdf>.
- . 2015. *The Importance of Cyber Threat Intelligence to a Strong Security Posture*. Ponemon Institute Research Report. www.webroot.com/shared/pdf/CyberThreatIntelligenceReport2015.pdf.
- Railton, Reanue. 2015. "When Business Process and Incident Response Collide: The Fine Tuning of the IR Program." Presentation at 27th Annual FIRST Conference, Berlin, Germany, June 16.
- Raustiala, Kal and David G. Victor. 2004. "The Regime Complex for Plant Genetic Resources." *International Organization* 58 (2): 277–309.
- Raymond, Mark and Gordon Smith, eds. 2014. *Organized Chaos: Reimagining the Internet*. Waterloo, ON: CIGI.

- Raymond, Mark and Laura DeNardis. 2015. "Multistakeholderism: Anatomy of an Inchoate Global Institution." *International Theory* 7 (3): 575–616.
- Rosenzweig, Paul. 2015. "The Administration's Cyber Proposals — Information Sharing." *Lawfare* (blog), January 16. www.lawfareblog.com/administrations-cyber-proposals-information-sharing.
- Ruefle, Robin, Audrey Dorofee, David Mundie, Allen D. Householder, Michael Murray and Samuel J. Perl. 2014. "Computer Security Incident Response Team Development and Education." *IEEE Security & Privacy* (September/October).
- Silver, Nate. 2015. *The Signal and the Noise: Why So Many Predictions Fail — But Some Don't*. New York: Penguin Books.
- Skierka, Isabel, Robert Morgus, Mirko Hohmann and Tim Maurer. 2015. "CSIRT Basics for Policy-Makers: The History, Types & Culture of Computer Security Incident Response Teams." Global Public Policy Institute Working Paper. GPPI, April 29. www.gppi.net/publications/global-internet-politics/article/csirt-basics-for-policy-makers/.
- Slaughter, Anne-Marie. 2006. "Networking Goes International: An Update." *Annual Review Law & Social Science* 2: 211–29.
- Struett, Michael J., Mark T. Nance and Diane Armstrong. 2013. "Navigating the Maritime Piracy Regime Complex: A Review of Multilateralism and International Organization." *Global Governance* 19 (1): 93–104.
- Trusted Introducer. 2015. "Services for Security and Incident Response Teams." Last modified May 5. www.trusted-introducer.org/.
- Verizon. 2015. *2015 Data Breach Investigations Report*. www.verizonenterprise.com/DBIR/2015/.
- Ward, Mark. 2014. "Cryptolocker Victims to Get Files Back for Free." BBC News, August 6. www.bbc.com/news/technology-28661463.
- Williams, Alex. 2013. "Target May be Liable for Up to 3.6 Billion from Credit Card Data Breach." *Tech Crunch*, December 23. <http://techcrunch.com/2013/12/23/target-may-be-liable-for-up-to-3-6-billion-from-credit-card-data-breach/>.
- Zetter, Kim. 2015. "Attackers Stole Certificate from FoxCon to Hack Kaspersky with DuQu 2.0." *Wired*, June 15. www.wired.com/2015/06/foxconn-hack-kaspersky-duqu-2/.
- Zheng, Denise E. and James A. Lewis. 2015. *Cyber Threat Information Sharing: Recommendations for Congress and the Administration*. March. Washington, DC: Center for Strategic & International Studies. http://csis.org/files/publication/150310_cyberthreatinfosharing.pdf.

ABOUT THE AUTHOR

Samantha Bradshaw is an expert on the high politics of Internet governance and cyber security technology. She joined CIGI as a research associate in October 2013 in the Global Security & Politics Program. She contributes to CIGI's work on Internet governance, and is a key member of a small team facilitating the Global Commission on Internet Governance. Samantha holds a joint Honours B.A. in political science and legal studies from the University of Waterloo and an M.A. in global governance from the Balsillie School of International Affairs.

CHAPTER NINE:
TOWARD A SOCIAL COMPACT FOR DIGITAL PRIVACY AND SECURITY
Statement by the Global Commission on Internet Governance

Copyright © 2015 by the Centre for International Governance Innovation
and the Royal Institute of International Affairs

SUMMARY

The Global Commission on Internet Governance (GCIG) was established in January 2014 to articulate and advance a strategic vision for the future of Internet governance. In recent deliberations, the Commission discussed the potential for a damaging erosion of trust in the absence of a broad social agreement on norms for digital privacy and security. The Commission considers that, for the Internet to remain a global engine of social and economic progress that reflects the world's cultural diversity, confidence must be restored in the Internet because trust is eroding. The Internet should be open, freely available to all, secure and safe. The Commission thus agrees that all stakeholders must collaborate together to adopt norms for responsible behaviour on the Internet. On the occasion of the April 2015 Global Conference on Cyberspace meeting in The Hague, the Commission calls on the global community to build a new social compact between citizens and their elected representatives, the judiciary, law enforcement and intelligence agencies, business, civil society and the Internet technical community, with the goal of restoring trust and enhancing confidence in the Internet.

It is now essential that governments, collaborating with all other stakeholders, take steps to build confidence that the right to privacy of all people is respected on the Internet. It is essential at the same time to ensure the rule of law is upheld. The two goals are not exclusive; indeed, they are mutually reinforcing. Individuals and businesses must be protected both from the misuse of the Internet by terrorists, cybercriminal groups and the overreach of governments and businesses that collect and use private data.

A social compact must be built on a shared commitment by all stakeholders in developed and less-developed countries to take concrete action in their own jurisdictions to build trust and confidence in the Internet. A commitment to the concept of collaborative security and to privacy must replace lengthy and over-politicized negotiations and conferences.

The following are the core elements that the Commission advocates in building the new social compact:

- Fundamental human rights, including privacy and personal data protection, must be protected online. Threats to these core human rights should be addressed by governments and other stakeholders acting both within their own jurisdiction and in cooperation.
- Interception of communications, collection, analysis and use of data over the Internet by law enforcement and government intelligence agencies should be for purposes that are openly specified in advance, authorized by law (including international human rights law) and consistent with the principles of necessity and proportionality. Purposes such as

gaining political advantage or exercising repression are not legitimate.

- In particular, laws should be publicly accessible, clear, precise, comprehensive and non-discriminatory, openly arrived at and transparent to individuals and businesses. Robust, independent mechanisms should be in place to ensure accountability and respect for rights. Abuses should be amenable to appropriate redress, with access to an effective remedy provided to individuals whose right to privacy has been violated by unlawful or arbitrary surveillance.
- Businesses or other organizations that transmit and store data using the Internet must assume greater responsibility to safeguard that data from illegal intrusion, damage or destruction. Users of paid or so-called “free services” provided on the Internet should know about, and have some choice over, the full range of commercial use on how their data will be deployed, without being excluded from the use of software or services customary for participation in the information age. Such businesses should also demonstrate accountability and provide redress in the case of a security breach.
- There is a need to reverse the erosion of trust in the Internet brought about by the non-transparent market in collecting, centralizing, integrating and analyzing enormous quantities of private information about individuals and enterprises — a kind of private surveillance in the service of “big data,” often under the guise of offering a free service.
- Consistent with the United Nations Universal Declaration of Human Rights, communications should be inherently considered private between the intended parties, regardless of communications technology. The role of government should be to strengthen the technology upon which the Internet depends and its use, not to weaken it.
- Governments should not create or require third parties to create “back doors” to access data that would have the effect of weakening the security of the Internet. Efforts by the Internet technical community to incorporate privacy-enhancing solutions in the standards and protocols of the Internet, including end-to-end encryption of data in transit and at rest, should be encouraged.
- Governments, working in collaboration with technologists, businesses and civil society, must help educate their publics in good cyber security practices. They must also collaborate to enhance the training and development of the software workforce globally, to encourage creation of more secure and stable networks around the world.
- The transborder nature of many significant forms of cyber intrusion curtails the ability of the target state

to interdict, investigate and prosecute the individuals or organizations responsible for that intrusion. States should coordinate responses and provide mutual assistance in order to curtail threats, to limit damage and to deter future attacks.

This statement provides the Commission's view of the issues at stake and describes in greater detail the core elements that are essential to achieving a social compact for digital privacy and security.

INTRODUCTION: THE OPPORTUNITIES AND RISKS EMERGING FROM THE INTERNET

In a short period of time, the Internet has become enmeshed in our daily lives. Now, people can exchange text, voice, images and data of all kinds — from anywhere in the world, instantly. We can create content, interact digitally, shop internationally with ease, exchange knowledge and ideas, and work together globally. The Internet, as a network of networks, is already capable of communicating and storing almost unimaginable volumes of data online, including data that can be associated with each of us individually and can be used for good or for ill.

In developed economies, the Internet has already delivered substantial social and economic benefits and is now an essential vehicle for innovation. For the developing world, the Internet can represent a powerful medium for social progress and economic growth, lifting millions of people out of poverty. For those struggling against repressive regimes, it represents a window into the wider world, a voice and a means to mobilize resistance and support. For those wishing to spread violent and hateful ideologies, it represents an unparalleled opportunity to try to radicalize

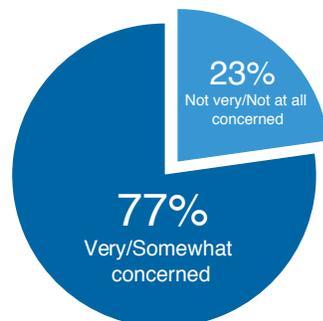
new audiences. For those seeking criminal gains, it represents a way of conducting traditional crimes on a larger scale and conducting new forms of Internet-enabled crime.

It is important to recognize that the communications and data of all of these actors are mixed together in the packet-switched networks and data clouds of the Internet. They all use the same fixed and, increasingly, mobile devices operating with the same Internet protocols. For the authorities charged with tracking down terrorists, countries that conduct espionage, cyber vandals and criminals of all kinds, the Internet provides a reservoir of information about their targets. But at the same time, the ability to access the intermingled data raises concerns over personal privacy and data protection.

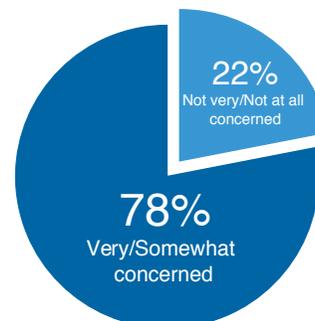
All developed economies now have multiple Internet dependencies. As the global reliance on the Internet rises, the vulnerability to disruption increases. Although Internet access is far from universal, by 2020 the number of Internet users is expected to reach five billion, with each user capable of interacting with any other. The largest portion of this further growth will be in the developing economies. The opportunities to collect, retain and use data for commercial profit, for harm and criminal gain, and for intelligence and security purposes, will increase commensurately. All stakeholders' capacity to protect fundamental human rights and to respond effectively will need to keep pace.

This shift in the availability of personal, commercial and public sector information, and the potential for access to infrastructure and control systems, represents a new source of vulnerability for society, magnified by the

Public Concern over Hacking of Personal Accounts



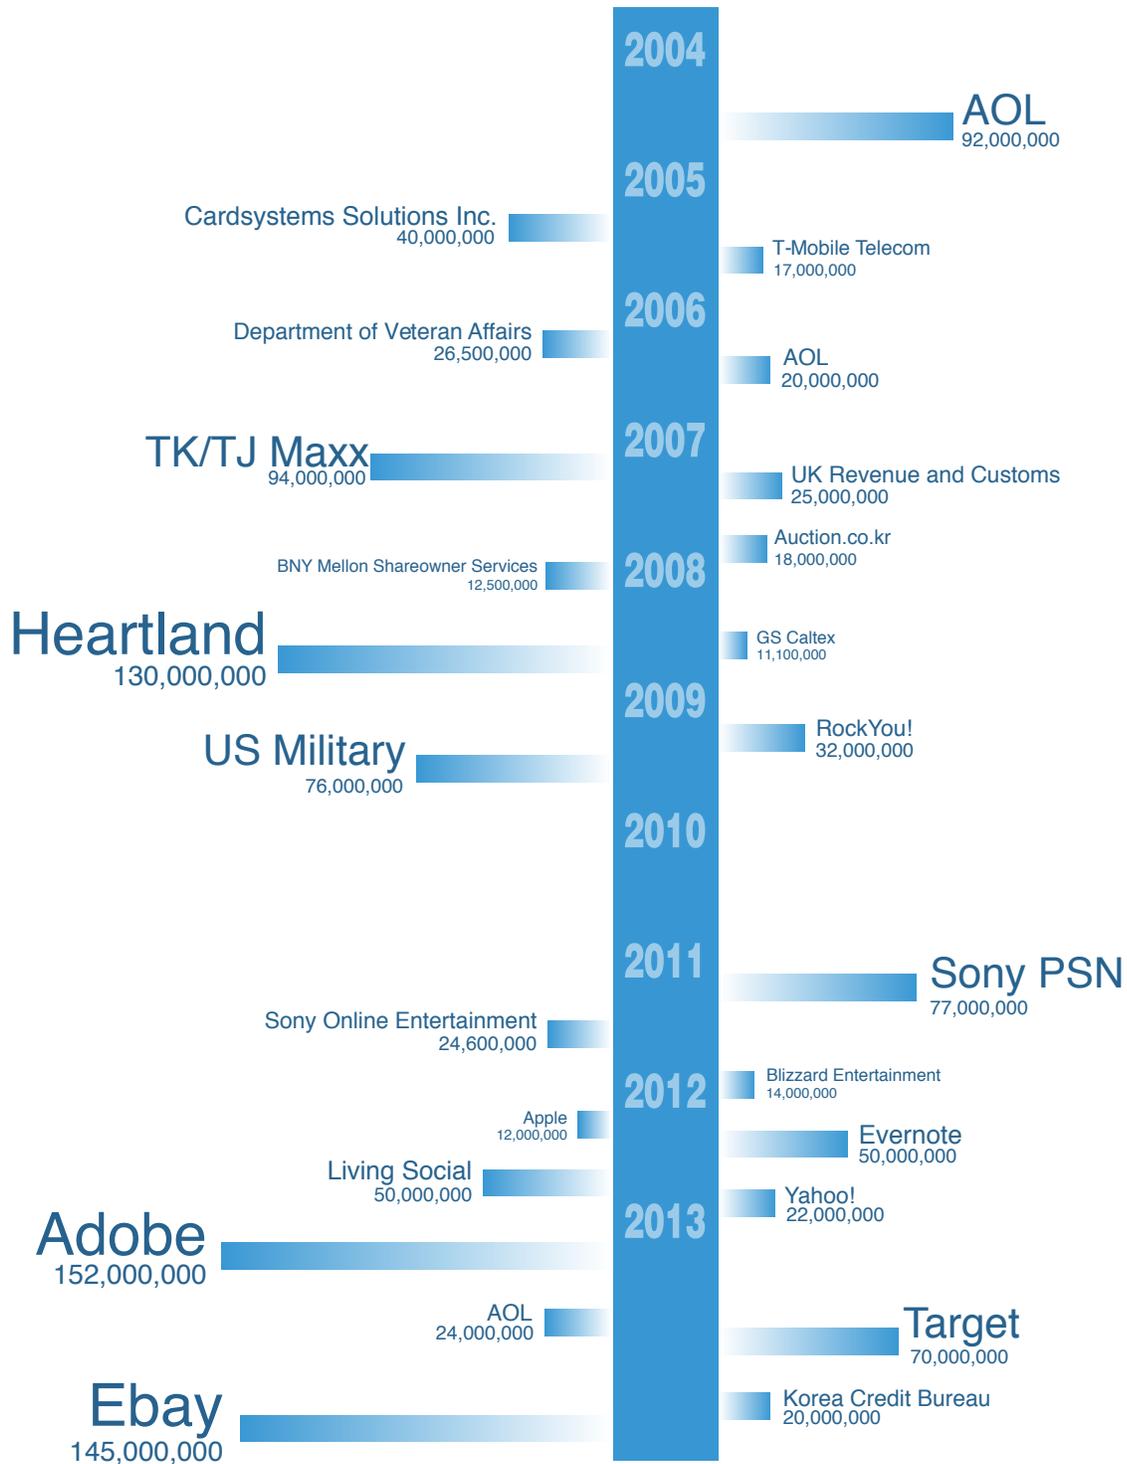
Public Concern over Personal Financial Cybercrime



Data source: CIGI-Ipsos Global Survey on Internet Security and Trust. Available at www.cigionline.org/internet-survey.

Note: The CIGI-Ipsos Global Survey was conducted between October 7, 2014, and November 12, 2014. Twenty-four countries were polled, including Australia, Brazil, Canada, China, Egypt, France, Germany, Great Britain, Hong Kong, India, Indonesia, Italy, Japan, Kenya, Mexico, Nigeria, Pakistan, Poland, South Africa, South Korea, Sweden, Tunisia, Turkey and the United States. In total, 23,326 Internet users were polled, aged 18–64 in Canada and the United States, aged 16–64 in every other country.

Data Breaches Affecting over 10 Million



Data source: Business Insider. Available at: www.businessinsider.com/data-breaches-infographic-2014-12.

growing use of mobile devices and wireless networks that offer additional ways for networks to be penetrated.

These dangers will be accentuated by the advent of the “Internet of Things” that is already starting to connect the key objects and instruments of daily life — our cars, our homes, our appliances, our clothing and much more. In the emerging world of the Internet of Things, everything we do, see, use or touch will leave electronic tracks, enlarging further both the potential commercial and social value of such data. It also will expand the opportunities provided for police and intelligence agencies to learn more about their suspects. Important questions still have to be addressed concerning the vulnerability of such connected systems and the privacy implications of allowing state and private-sector actors to have access to and to share the big data that they will generate. Similarly, there will be a need to clarify that whatever access there is must have a legal basis.

INDIVIDUALS, BUSINESSES AND GOVERNMENTS FACE NEW CHALLENGES

This data revolution has significant and complex negative implications for three sets of actors: individuals, businesses and governments.

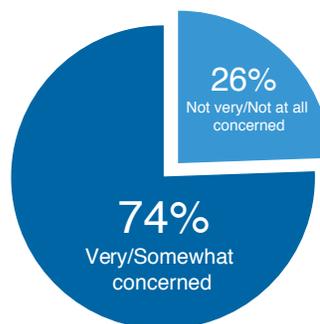
A number of surveys indicate that, for individual and corporate users of the Internet, the primary concern is to have adequate assurance of the security of their information against misuse: the cybercrime, vandalism, theft and even terrorist acts that the Internet enables. Not all individuals understand the full scope of what they have placed online deliberately or what information has been captured and stored by others as they go about their daily activities. Nor do most individuals know to what commercial use their data are deployed.

Third parties who have access to data have the potential to monitor, obtain and put to use enormous quantities of private information about individuals and businesses, their communications, their plans, their locations and behaviour, even their shopping, viewing and reading habits. These developments and increasing awareness of them pose a substantial challenge to safety and security, to privacy rights and to citizens’ trust in the Internet, which has steadily been eroding. Therefore, these developments are also a substantial threat to the social and economic value of the Internet.

Today, some companies exceed governments in their capacity to collect, store in centralized repositories, integrate, analyze and make use of personal data. These companies are increasingly attractive targets for cyber intrusion, and susceptible to efforts to jeopardize the confidentiality, availability and integrity of these large data pools. These companies have to demonstrate to their users a high level of respect for, and protection of, the security and privacy of their information. At the same time, companies must exhibit corporate social responsibility in responding to government requests for access to their users’ data. They also must contend with increasing requests for access to data from law enforcement overseas due to the transborder nature of many activities taking place on the Internet.

Many companies operating on the Internet also are building their businesses on the use and sale of the data they gather. Often the data are accessed in exchange for providing a free service to their users. Data collected from customers are often used for purposes not explicitly revealed to those who provide the data, and used without their permission. On one hand, this is fuelling data analytics to the benefit of innovation. On the other, it raises concerns about the respect for users’ privacy. There is a rising call for regulators, or for the industry itself, to establish standards for transparency and accountability mechanisms to increase confidence in the marketplace.

Internet Users’ Concern over Private Company Monitoring of Online Activity and Sale of User-generated Data



Data source: CIGI-Ipsos Global Survey on Internet Security and Trust. Available at www.cigionline.org/internet-survey.

Governments have the responsibility to pursue Internet policies that are consistent with fundamental human rights and the rule of law, and that promote economic well-being. At the same time, they have a duty to address threats from both state and so-called “non-state actors” such as dictators, insurgents, terrorists and other criminals of all kinds. As data and communications of all types moved from traditional telephone and radio technologies to Internet-based transmission, the opportunities for intelligence agencies to monitor such targets by intercepting and exploiting digital data increased. Yet it is difficult for law enforcement officials to interdict and prosecute transnational criminal activity without having assistance from secret intelligence agencies and their powerful tools of digital intelligence gathering. For example, the pattern and content of messages sent between al-Qaeda, Boko Haram, ISIL (Islamic State of Iraq and the Levant) or other terrorist operatives, and those between members of transnational criminal organizations, would be a high priority for interception by the intelligence and law enforcement agencies of many nations. Cooperation may be required to share specialized resources, because a great deal of criminal and socially damaging activity takes place in the deep recesses of the Internet, including the so-called “dark Web.” Oversight is required to assure citizens that their rights are not infringed upon in the pursuit of a range of bad actors.

Government activities themselves are vulnerable to terrorists and cybercriminals through the Internet. Many governments are seeking to work with businesses to improve national cyber security to counter the risks of cybercrime, disruption and destruction, especially of critical national infrastructure. These increased risks underscore the importance of governments monitoring threats and attacks online. Nevertheless, some governments are conducting both targeted and mass surveillance in ways that have a chilling effect on fundamental human rights and, in particular, freedom of expression and legitimate dissent and protest, and that threaten the realization of the Internet’s economic and social benefits.

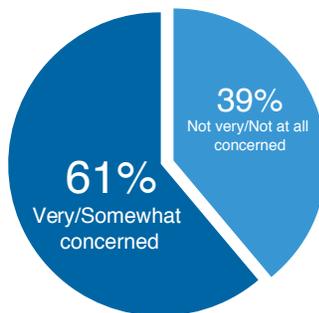
NATIONAL AND INTERNATIONAL RESPONSES

The speed of these contradictory developments in the use of the Internet has left policy lagging behind. Governments struggle to know how to manage the harms the Internet facilitates while preserving its power for good.

At a domestic level, responding to pressure from privacy and civil liberties organizations, in several nations a debate has started about the nature, capacity and legal framework of their digital intelligence activities. Some Internet and telecommunications companies now publish transparency reports about the demands governments place on them. Some nations already have comprehensive legislation to regulate intrusive digital intelligence powers; others do not. Some have parliamentary or judicial oversight (or both) of such activity while some do not have either. Personal data protection regulations are mostly not yet suited to the complexity of the digital age — for example, by not adequately regulating the extensive secondary use of personal data or ensuring the transparency of exceptions to privacy for sovereignty and national security purposes. The military utility of offensive cyber operations and intelligence attacks is increasingly recognized, as are the dangers posed by advanced malware and software flaws.

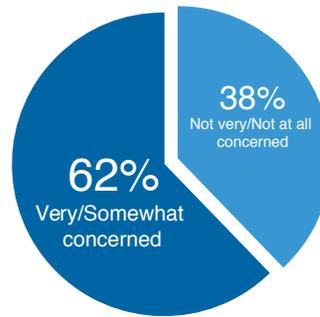
At the international level, all states have subscribed to the UN Universal Declaration on Human Rights, and almost all states have ratified the UN International Covenant on Civil and Political Rights, both of which enshrine the right to privacy in international human rights law. Additionally, some groups of states have usefully developed the right to privacy further, such as in the Convention on Human Rights from the Council of Europe and by implementing the judgments of the European Court of Human Rights. Furthermore, both the NETmundial outcome document and the two recently adopted resolutions from the UN General Assembly on the Right to Privacy in the Digital Age affirmed that the same rights that people have offline

Public Concern over Domestic State Surveillance



Data source: CIGI-Ipsos Global Survey on Internet Security and Trust. Available at www.cigionline.org/internet-survey.

Public Concern with Foreign Government Surveillance



Data source: CIGI-Ipsos Global Survey on Internet Security and Trust. Available at www.cigionline.org/internet-survey.

must also be protected online, including the right to privacy.

The obligation of states to protect and promote rights to privacy and freedom of expression are not optional. Even if they are not absolute rights, limitations to these rights, even those based on national security concerns, must be prescribed by law, guaranteeing that exceptions are both necessary and proportionate. Governments should guarantee the same human rights protection to all individuals within their borders. Clearly, any interference with the right to privacy should not be arbitrary or unlawful, bearing in mind what is reasonable to the pursuance of legitimate aims. The Organisation for Economic Co-operation and Development (OECD) Privacy Guidelines state that exceptions to its principles, including those relating to national sovereignty, national security and public policy (*ordre public*), should be as few as possible, and made known to the public. The 2013 International Principles on the Application of Human Rights to Communications Surveillance, developed at the initiative of civil society, are an important reference regarding how international human rights law should apply in the current digital environment. States are called to comply with the following principles: legality, legitimate aim, necessity, adequacy, proportionality, competent judicial authority, due process, user notification, transparency, public oversight, integrity of communications and systems, safeguards for international cooperation, safeguards against illegitimate access and the right to effective remedy.

Formal and informal efforts such as these are early steps in the emergence of a new social compact for the digital age.

CORE ELEMENTS OF A SOCIAL COMPACT FOR A DIGITAL SOCIETY

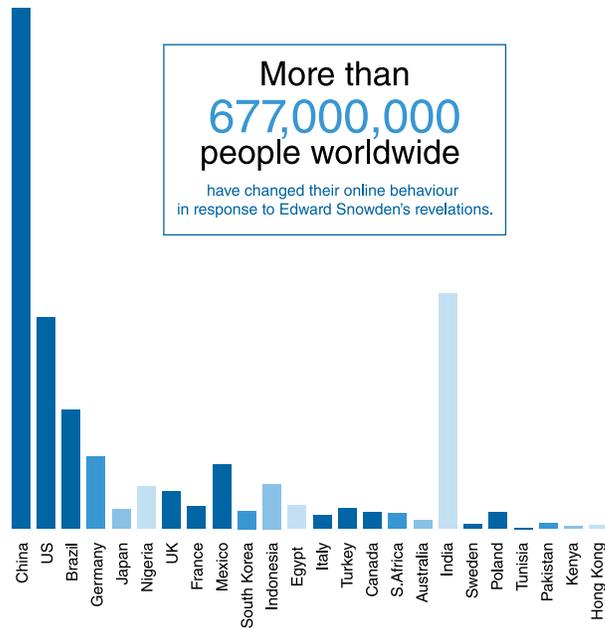
There must be a mutual understanding between citizens and their state that the state takes responsibility to keep its citizens safe and secure under the law while, in turn, citizens

agree to empower the authorities to carry out that mission, under a clear, accessible legal framework that includes sufficient safeguards and checks and balances against abuses. Business must be assured that the state respects the confidentiality of its data and companies must, in turn, provide their customers the assurance that their data is not misused. There is an urgent need to achieve consensus on a social compact for the digital age in all countries. Just how urgent is shown by current levels of concern over allegations of intrusive state-sponsored activities, ranging from weakening of encryption to large-scale criminal activity, to digital surveillance, to misuse of personal data and even to damaging cyber attacks and disruption.

In an environment of rapidly changing technologies and social attitudes, a normative approach would be a practical starting point for such an effort. Key elements of a social compact for the digital age will necessarily take different institutional and legal forms in different societies and cultures. Nevertheless, a global social compact should be informed by a number of core elements:

- Fundamental human rights, including privacy and personal data protection, must be protected online. Threats to these core human rights should be addressed by governments and other stakeholders acting both within their own jurisdiction and in cooperation.
- Interception of communications, collection, analysis and use of data over the Internet by law enforcement and government intelligence agencies should be for purposes that are openly specified in advance, authorized by law (including international human rights law) and consistent with the principles of necessity and proportionality. Purposes such as gaining political advantage or exercising repression are not legitimate.
- In particular, laws should be publicly accessible, clear, precise, comprehensive and non-discriminatory, openly arrived at and transparent to individuals and businesses. Robust, independent mechanisms should

Online Behavioural Change in Response to Edward Snowden's Revelations



Data source: CIGI-Ipsos Survey on Internet Security and Trust, available at www.cigionline.org/internet-survey and World Bank Indicators on Population and Internet Penetration Rates, available at <http://data.worldbank.org/indicator>.

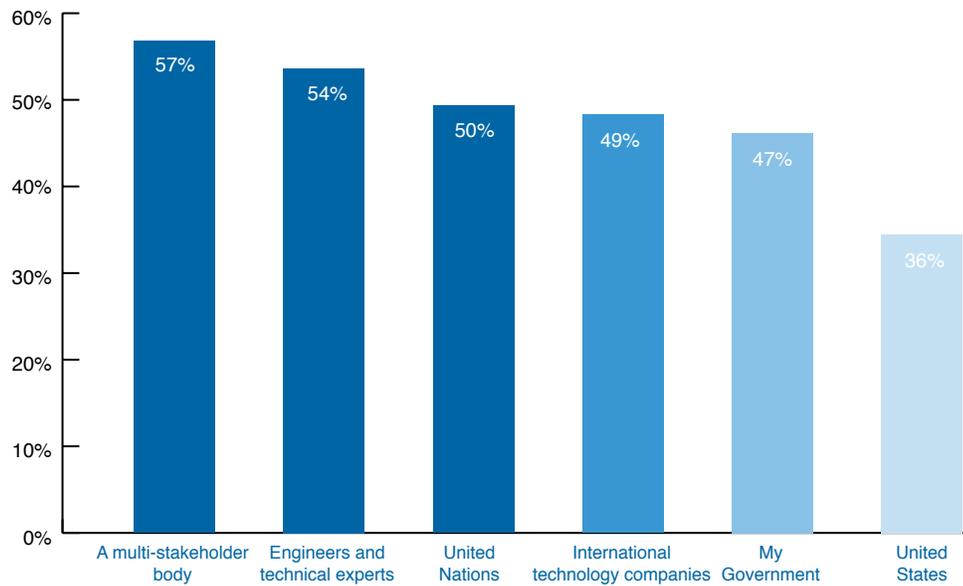
be in place to ensure accountability and respect for rights. Abuses should be amenable to appropriate redress, with access to an effective remedy provided to individuals whose right to privacy has been violated by unlawful or arbitrary surveillance.

- Businesses or other organizations that transmit and store data using the Internet must assume greater responsibility to safeguard that data from illegal intrusion, damage or destruction. Users of paid or so-called “free services” provided on the Internet should know about, and have some choice over, the full range of commercial use on how their data will be deployed, without being excluded from the use of software or services customary for participation in the information age. Such businesses should also demonstrate accountability and provide redress in the case of a security breach.
- There is a need to reverse the erosion of trust in the Internet brought about by the non-transparent market in collecting, centralizing, integrating and analyzing enormous quantities of private information about individuals and enterprises — a kind of private surveillance in the service of “big data,” often under the guise of offering a free service.
- Consistent with the United Nations Universal Declaration of Human Rights, communications should be inherently considered private between the intended parties, regardless of communications technology. The role of government should be to

strengthen the technology upon which the Internet depends and its use, not to weaken it.

- Governments should not create or require third parties to create “back doors” to access data that would have the effect of weakening the security of the Internet. Efforts by the Internet technical community to incorporate privacy-enhancing solutions in the standards and protocols of the Internet, including end-to-end encryption of data in transit and at rest, should be encouraged.
- Governments, working in collaboration with technologists, businesses and civil society, must help educate their publics in good cyber security practices. They must also collaborate to enhance the training and development of the software workforce globally, to encourage creation of more secure and stable networks around the world.
- The transborder nature of many significant forms of cyber intrusion curtails the ability of the target state to interdict, investigate and prosecute the individuals or organizations responsible for that intrusion. States should coordinate responses and provide mutual assistance in order to curtail threats, to limit damage and to deter future attacks.

The Public's Preference for Multi-stakeholder Governance



Data source: CIGI-Ipsos Global Survey on Internet Security and Trust. Available at www.cigionline.org/internet-survey.

MOVING TOWARD A SOCIAL COMPACT FOR A DIGITAL SOCIETY

The social compact for a digital society will require a very high level of agreement among governments, private corporations, individuals and the technical community. Governments can provide leadership, but cannot alone define the content of the social compact. Achieving agreement and acceptance will necessitate the engagement of all stakeholders in the Internet ecosystem. At first, it is unlikely that a universal social compact suitable to all circumstances could, or even should, be the immediate goal. The Internet is used and valued across all cultures and all borders. Significant changes of attitude can sometimes evolve more quickly and more flexibly than could be possible through negotiated treaties or international legal instruments. In the fullness of time, national approaches may gain recognition as good international practices, and may eventually acquire the status of customary international law. But that is many years away, and the speed of technological change argues for flexibility and innovative solutions. The area of secret intelligence is especially difficult to regulate since there is little international law governing it, but even that largely secret domain ought not to be free of ethical and legal considerations.

The social compact will contribute to building a new kind of “collaborative privacy and security.” The term highlights a fundamental truth about the Internet: every part of the Internet ecosystem affects every other part. Thus, the new social compact is not about “balancing” human rights and privacy against states’ interests or against commercial rights. It is about ensuring that a framework exists

where each actor has the responsibility to act not only in their own interest, but also in the interest of the Internet ecosystem as a whole. By definition, the process should result in outcomes that are win-win rather than zero-sum games. Effective security, successful business models and human rights are mutually reinforcing in the long run. All interests must recognize and act on their responsibility for security and privacy on the Internet in collaboration with all others, or no one is successful.

In the end, it is in the interest of all stakeholders that the Internet remains trusted as a common global resource: open, affordable, unfettered and available to all as a safe medium for further innovation. Government, business and civil society must work together toward that aim.

CONCLUSION

These recommendations are put forward by the Global Commission on Internet Governance to encourage a strong consensus among all stakeholders that the benefits of the Internet for humankind must not be put at risk, whether by disproportionate state behaviour in cyberspace, by criminal activity or by business activity undermining assurance in the confidentiality, integrity and availability of information on the Internet. Advancing a new normative framework, which accounts for the dynamic interplay between national security interests and the needs of law enforcement, while preserving the economic and social value of the Internet, is an important first step to achieving long-term digital trust. The Commission is committed to building on this statement by continuing its program of research and publication, undertaken in collaboration with partners from all sectors.

ACKNOWLEDGEMENTS

The Commissioners wish to thank Bill Graham, CIGI senior fellow, and Aaron Shull, CIGI fellow, for their assistance in drafting this statement; Eric Jardine, CIGI research fellow, for preparing the charts and figures that appear in the document; and Samantha Bradshaw, CIGI research associate, for her research assistance.

FURTHER READING

In developing this statement, the Commissioners drew upon a number of publications that outline the issues of trust, privacy and security. A partial list of the works consulted and others recommended for further research can be found at www.ourinternet.org. This is a representative list only, and is not intended to be complete or comprehensive.

Barnes, R. et al. 2015. "Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement." Network Working Group, Internet Engineering Task Force. www.ietf.org/id/draft-iab-privsec-confidentiality-threat-04.txt.

Bartlett, J. and Alex Krasodowski-Jones. 2015. "Online Anonymity, Islamic State and Surveillance." Demos. www.demos.co.uk/files/Islamic_State_and_Encryption.pdf?1426713922.

Bildt, Carl. 2013. Speech by Foreign Minister Carl Bildt at Seoul Conference on Cyberspace 2013, Seoul, October 17. Utrikesdepartementet, Sweden. www.government.se/sb/d/7956/a/226592.

Chertoff, Michael and Toby Simon. 2015. *The Impact of the Dark Web on Internet Governance and Cyber Security*. Global Commission on Internet Governance Paper No. 6. Waterloo: CIGI and Chatham House. https://ourinternet-files.s3.amazonaws.com/publications/GCIG_Paper_No6.pdf.

CIGI/IPSOS. 2014. "Global Survey on Internet Security and Trust." Waterloo, ON: CIGI. www.cigionline.org/internet-survey.

Daigle, Leslie. 2015. *On the Nature of the Internet*. Global Commission on Internet Governance Paper No. 9. Waterloo, ON: CIGI and Chatham House. <https://ourinternet.org/#publications/on-the-naure-of-the-internet>.

Electronic Frontier Foundation et al. 2014. "The International Principles on the Application of Human Rights to Communications Surveillance." May. <https://en.necessaryandproportionate.org/>.

Internet Society. 2015a. "Internet Society Approach to Cyber Security Policy." Internet Society. January 22. www.internetsociety.org/news/internet-society-approach-cyber-security-policy.

———. 2015b. "Understanding Security and Resilience of the Internet." Internet Society. www.internetsociety.org/sites/default/files/bp-securityandresilience-20130711.pdf.

La Rue, Frank. 2013. *Report of the Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression, Frank La Rue*. United Nations General Assembly: Human Rights Council. www.ohchr.org/Documents/HRBodies/HRCouncil/RegularSession/Session23/A.HRC.23.40_EN.pdf.

NETmundial. 2014. "NETmundial Draft Outcome Statement." Global Multistakeholder Meeting on the Future of Internet Governance. http://document.netmundial.br/net-content/uploads/2014/04/NETmundial-draft-outcome-document_April_14.pdf.

Nye, Joseph S., Jr. 2014. *The Regime Complex for Managing Global Cyber Activities*. Global Commission on Internet Governance Paper No. 1. Waterloo, ON: CIGI and Chatham House. www.cigionline.org/sites/default/files/gcig_paper_no1.pdf.

OECD. 2013. *Recommendation of the Council concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data (The Privacy Guidelines)*. OECD. www.oecd.org/sti/ieconomy/2013-oecd-privacy-guidelines.pdf.

OECD. 2014. *OECD Principles for Internet Policy Making*. OECD. www.oecd.org/sti/ieconomy/oecd-principles-for-internet-policy-making.pdf.

Office of the High Commissioner for Human Rights. 2004. "Nature of the General Legal Obligation on States Parties to the Covenant." General Comment 31, Office of the High Commissioner for Human Rights. www1.umn.edu/humanrts/gencomm/hrcom31.html.

Omand, David. 2015. *Understanding Digital Intelligence and the Norms That Might Govern It*. Global Commission on Internet Governance Paper No. 8. Waterloo, ON: CIGI and Chatham House. https://ourinternet-files.s3.amazonaws.com/publications/gcig_paper_no8.pdf.

Scheinin, Martin. 2010. *Report of the Special Rapporteur on the Promotion and Protection of Human Rights and Fundamental Freedoms while Countering Terrorism, Martin Scheinin*. United Nations General Assembly: Human Rights Council. <https://fas.org/irp/eprint/unhrc.pdf>.

- United Nations. 2011. "Guiding Principles on Business and Human Rights: Implementing the United Nations 'Protect, Respect and Remedy' Framework." United Nations Human Rights Council Resolution 17/4. www.ohchr.org/Documents/Publications/GuidingPrinciplesBusinessHR_EN.pdf.
- . 2012. "Brazil and Germany: Draft Resolution: The Right to Privacy in the Digital Age." Draft Resolution: Sixty-Eighth Session, Third Committee, UN General Assembly. www.hrw.org/sites/default/files/related_material/UNGA_upload_0.pdf.
- . 2015. "The Right to Privacy in the Digital Age." Resolution Adopted by the UN General Assembly on December 18, 2014. www.un.org/en/ga/search/view_doc.asp?symbol=A/RES/69/166.
- Verhulst, Stefaan G. et al. 2014. *Innovations in Global Governance: Toward a Distributed Governance Ecosystem*. Global Commission on Internet Governance Paper No. 5. Waterloo, ON: CIGI and Chatham House. https://ourinternet-files.s3.amazonaws.com/publications/gcig_paper_no5.pdf.
- Weber, Rolf. 2014. *Legal Interoperability as a Tool for Combatting Fragmentation*. Global Commission on Internet Governance Paper No. 4. Waterloo, ON: CIGI and Chatham House. https://ourinternet-files.s3.amazonaws.com/publications/gcig_paper_no4.pdf.

GLOBAL COMMISSIONERS

Carl Bildt

Chair of the Global Commission on Internet Governance

Gordon Smith

Deputy Chair of the Global Commission on Internet Governance

Fen Osler Hampson

Co-Director of the Global Commission on Internet Governance

Patricia Lewis

Co-Director of the Global Commission on Internet Governance

Laura DeNardis

Director of Research of the Global Commission on Internet Governance

Sultan Sooud Al Qassemi

Dominic Barton

Pablo Bello

Pascal Cagni

Moez Chakchouk

Dae-Whan Chang

Michael Chertoff

Dian Triansyah Djani

Anriette Esterhuysen

Hartmut Glaser

Dorothy Gordon

Angel Gurría

Dame Wendy Hall

Melissa Hathaway

Mathias Müller von Blumencron

Beth Simone Noveck

Joseph S. Nye, Jr.

Sir David Omand

Nii Quaynor

Latha Reddy

Marietje Schaake

Tobby Simon

Michael Spence

Paul Twomey

Pindar Wong

ABOUT CIGI

We are the Centre for International Governance Innovation: an independent, non-partisan think tank with an objective and uniquely global perspective. Our research, opinions and public voice make a difference in today's world by bringing clarity and innovative thinking to global policy making. By working across disciplines and in partnership with the best peers and experts, we are the benchmark for influential research and trusted analysis.

Our research programs focus on governance of the global economy, global security and politics, and international law in collaboration with a range of strategic partners and support from the Government of Canada, the Government of Ontario, as well as founder Jim Balsillie.

Au Centre pour l'innovation dans la gouvernance internationale (CIGI), nous formons un groupe de réflexion indépendant et non partisan qui formule des points de vue objectifs dont la portée est notamment mondiale. Nos recherches, nos avis et l'opinion publique ont des effets réels sur le monde d'aujourd'hui en apportant autant de la clarté qu'une réflexion novatrice dans l'élaboration des politiques à l'échelle internationale. En raison des travaux accomplis en collaboration et en partenariat avec des pairs et des spécialistes interdisciplinaires des plus compétents, nous sommes devenus une référence grâce à l'influence de nos recherches et à la fiabilité de nos analyses.

Nos programmes de recherche ont trait à la gouvernance dans les domaines suivants : l'économie mondiale, la sécurité et les politiques mondiales, et le droit international, et nous les exécutons avec la collaboration de nombreux partenaires stratégiques et le soutien des gouvernements du Canada et de l'Ontario ainsi que du fondateur du CIGI, Jim Balsillie.

For more information, please visit www.cigionline.org.

ABOUT CHATHAM HOUSE

Chatham House, the Royal Institute of International Affairs, is based in London. Chatham House's mission is to be a world-leading source of independent analysis, informed debate and influential ideas on how to build a prosperous and secure world for all. The institute: engages governments, the private sector, civil society and its members in open debates and confidential discussions about significant developments in international affairs; produces independent and rigorous analysis of critical global, regional and country-specific challenges and opportunities; and offers new ideas to decision-makers and -shapers on how these could best be tackled from the near- to the long-term. For more information, please visit: www.chathamhouse.org.

CIGI MASTHEAD

Executive

President	Rohinton P. Medhora
Director of Finance	Shelley Boettger
Director of the International Law Research Program	Oonagh Fitzgerald
Director of the Global Security & Politics Program	Fen Osler Hampson
Director of Human Resources	Susan Hirst
Director of the Global Economy Program	Domenico Lombardi
Chief Operating Officer and General Counsel	Aaron Shull
Director of Communications and Digital Media	Spencer Tripp

Publications

Publisher	Carol Bonnett
Senior Publications Editor	Jennifer Goyder
Publications Editor	Patricia Holmes
Publications Editor	Nicole Langlois
Publications Editor	Sharon McCartney
Publications Editor	Lynn Schellenberg
Graphic Designer	Melodie Wakefield

For publications enquiries, please contact publications@cigionline.org.

Communications

For media enquiries, please contact communications@cigionline.org.



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

CHATHAM HOUSE

The Royal Institute of
International Affairs

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

