

Future of Digital Economy and Society System Initiative

Cyber Resilience Playbook for Public- Private Collaboration

In collaboration with The Boston Consulting Group

January 2018



Contents

Preface	3
1. Introduction	5
2. Using the Playbook for Public-Private Collaboration	6
3. Reference architecture for public-private collaboration	8
4. Policy models	11
4.1 Zero-days	11
4.2 Vulnerability liability	14
4.3 Attribution	19
4.4 Research, data, and intelligence sharing	22
4.5 Botnet disruption	26
4.6 Monitoring	30
4.7 Assigning national information security roles	33
4.8 Encryption	37
4.9 Cross-border data flows	41
4.10 Notification requirements	44
4.11 Duty of assistance	47
4.12 Active defence	51
4.13 Liability thresholds	54
4.14 Cyberinsurance	57
5. The future of cyber resilience	60
Appendix: Normative trade-offs framework	65
Acknowledgements	65
Endnotes	70

Preface

The World Economic Forum System Initiative on Shaping the Future of Digital Economy and Society represents a global platform for multistakeholder coalitions from across the world to collaborate and accelerate progress against shared digital economy goals and to shape a digital future that is sustainable, inclusive and trustworthy. This future requires leaders to build and foster institutions that meet the challenges of cybersecurity and help to mitigate cyber-risk across our shared networks.

Cyber-risk is and will continue to be one of the most pressing challenges accompanying the Fourth Industrial Revolution. Leaders across the public and private sectors appreciate that mitigating this risk requires continued collaboration. The Forum has led discussions on this topic since 2012 and this year will be inaugurating the Global Cyber Centre as a platform to continue advancing cyber resilience.

Collaboration is often difficult in the sphere of cybersecurity. Not only has technological innovation begun to implicate core societal values, the interdisciplinary dialogue required to collaborate and make progress often spans across many competencies, from the technical to the ethical.

To help frame discussion for leaders in both the public and private sectors, as part of the World Economic Forum System Initiative on Shaping the Future of Digital Economy and Society, the Forum has partnered with The Boston Consulting Group to develop a baseline framework to serve as a springboard for cooperation and shared understanding in cybersecurity policy-making. This report is the result of extensive collaboration, debate, consultation, and iteration to distil complex and nuanced issues in cybersecurity to their irreducible core.

The Forum would like to thank The Boston Consulting Group for its leadership, the Steering Committee and the Expert Working Group for their contributions, as well as the numerous leaders in cybersecurity who patiently helped shape our efforts this past year. This was an effort of multiple communities across industries and sectors and we are sincerely grateful for each of our partners' and contributors' dedication to this vital work.

We hope this document can begin fruitful collaboration to help advance our shared cyber resilience.



Cheryl Martin
Member of the
Managing Board



Rick Samans
Member of the
Managing Board

“

Governments of the Industrial World, you weary giants of flesh and steel, I come from Cyberspace, the new home of Mind. On behalf of the future, I ask you of the past to leave us alone. You are not welcome among us. You have no sovereignty where we gather.

John Perry Barlow, "A Declaration of the Independence of Cyberspace", Davos, 1996¹

”

“

Every week there are reports of newly discovered security problems in all kinds of software, from individual applications and services to Windows, Linux, Unix and other platforms. We have done a great job of having teams work around the clock to deliver security fixes for any problems that arise. Our responsiveness has been unmatched — but as an industry leader we can and must do better... We need to make it automatic for customers to get the benefits of these fixes. Eventually, our software should be so fundamentally secure that customers never even worry about it.

Bill Gates, "Trustworthy Computing", 2002

”

“

Like in the real world, freedom and order are both necessary in cyberspace. Freedom is what order is meant for and order is the guarantee for freedom. We should respect internet users' rights to exchange their ideas and express their minds, and we should also build a good order in cyberspace in accordance with law as it will help protect the legitimate rights and interests of all internet users. Cyberspace is not a place beyond the rule of law. Cyberspace is virtual, but players in cyberspace are real.

Xi Jinping, "At the Opening Ceremony of the Second World Internet Conference", 2015

”

1. Introduction

States have an obligation to provide security for their citizens. The increasingly networked, digitized and connected world has enlarged and complicated that obligation. These changes have also created new obligations, shared among a variety of actors, from states to corporations to civil society and individuals.

To meet this rapidly expanding obligation, leaders have taken a variety of approaches to securing their digital domains. These policies are shaped by their experience with the networked world and unique national objectives and vulnerabilities. For all their differences, however, these policy approaches to assuring security share a significant commonality: success depends on collaboration between the public and private sectors.

However, effective collaboration is uniquely difficult in the domain of cybersecurity. Cyberthreats are complex, with an ever-expanding and exposed surface for malicious actors to exploit. Each new innovation brings with it new and sometimes unexpected vulnerabilities. That complexity is compounded by the speed and ease with which threats materialize in

the digital domain — no expensive “Manhattan Project” style effort is necessary to weaponize computer science. Additionally, the first line of security here is rarely the government. Rather, the first line of security is comprised of the firms and organizations developing this increasingly networked, digitized and connected space.

Public-private collaboration is almost always difficult because of the complexity underlying the interplay between the roles, responsibilities and obligations that the public and private sectors have vis-à-vis each other and the citizens who rely on them. The difficulties of public-private collaboration are magnified when a topic, such as security, is deeply connected to notions of sovereignty: multinational businesses and customers walk a tightrope between potentially contradictory national obligations.

In the case of cybersecurity, that tension is further strained by the decidedly personal nature of securing bits and pieces of an increasing portion of people’s lives. The relationship and — at times trade-off — between security and other values magnifies the need to be inclusive in representing and negotiating between different interests and principles.

Despite these challenges, advancing cyber resilience requires the public and private sectors to collaborate in new and innovative ways. This Playbook is recommended for use by the public and private sectors, together, as a tool to facilitate discussions on building the institutions, frameworks, policies, norms and processes necessary to support collaboration in this vital space.

2. Using the Playbook for Public-Private Collaboration

The Playbook is intended to guide intra-state public-private collaboration on cybersecurity policy. This Playbook contains two distinct sections in service of that mission: the Reference architecture for public-private collaboration and the Cyber policy models.

Policy-makers and senior executives should begin by reviewing the Reference architecture for public-private collaboration for an overview of cybersecurity policy issues. After reviewing the Reference architecture, it is advisable to turn to a given policy question of interest and review the policy models, which frame each policy question.

Reference architecture for public-private collaboration

While leaders are accustomed to debating cybersecurity policy topics in isolation, there is seldom reflection on whether the sum of the parts of cybersecurity policy crafted on a day-to-day basis amounts to a coherent whole. It is easy to get lost in the particulars of any specific policy and neglect the unintended consequences of a given policy position on the broader edifice of cybersecurity policy. To help facilitate that discussion, the Reference architecture documents the key policy topics as well as some of the interdependencies that policy-makers should keep in mind (e.g. how threat intelligence sharing impacts the formation and disruption of botnets).

Policy model

Each policy model provides a brief reference for a specific topic, helping leaders in the public and private sectors to develop a baseline understanding of the key issues. In particular, these models provide an analytical framework for approaching policy questions, and document the risks and trade-offs associated with each policy, importantly including the normative trade-offs as well. Where appropriate, these models include case studies that illustrate a key concept surfaced by the topic.

The intent of describing trade-offs is not to advance specific policy positions which “should” be taken. Rather, it is to frame the different choices that “could” be made, with the goal of encouraging clear-eyed discussion and debate.


This document will also not enumerate how to operationally implement a specific policy. Rather, the aim is to abstract away from any individual country’s context to provide a common language to discuss cybersecurity policy generally. In practice, implementation will vary by national context: every country has unique latent capabilities, risks, and normative values.

Connecting policy to values

Throughout this discussion of different policy models, on topics ranging from zero-days to attribution, this document will attempt to connect policy positions to the norms and values that those positions prioritize or embody. The intent is to discourage polarization in security dialogue and move beyond the rhetorical simplicity of prioritizing one value over all others (e.g. “privacy cannot exist without security”) or a false-choice narrative that freezes action-oriented debate into prolonged indecision.

In connecting norms and values to policy positions, this document encourages all actors to move past absolute and rigid positions towards more nuanced discussions. To encourage these discussions, the Playbook discusses the implications of policy choices on five key values: security, privacy, economic value, accountability and fairness.

These values were selected on the basis of the judgement of our Working Group, given its experience in the security ecosystem after considering more than 20 different values ranging from interoperability to social cohesion. For a detailed overview of these values were considered, from policy evaluation to normative judgement, please see “Normative trade-offs framework” in the appendix.



3. Reference architecture for public-private collaboration

In contextualizing cybersecurity policy, 14 key policy topics dot the policy landscape.

1. Research, data and intelligence sharing
What is the government's role in sharing and promoting the dissemination of threat intelligence?

2. Zero-days
To what extent should the government be involved in the research, development and purchase of zero-day vulnerabilities and exploits?

To what extent should government share these vulnerabilities with the private sector?

3. Vulnerability liability
Who is liable for securing a vulnerability?

How should that liability shift if/when products transition to end-of-life?

4. Attribution
How should government engage with the private sector when the private sector publicly alleges that a particular actor is responsible for a given attack?

5. Botnet disruption
What should be done to prevent the proliferation of botnets?

How should existing botnets be researched and studied?

How should actors throughout the ecosystem disrupt botnets?

6. Monitoring
What should non-users be able to monitor to promote security and other valid national interests?

7. Assigning national information security roles
Which entities and organizations should be responsible for fulfilling different national information security roles?

8. Encryption
Who should be able to access sensitive data and communications?

9. Cross-border data flows
What are the security and non-security implications of countries exerting control over data?

10. Notification requirements
When should companies be required to notify relevant stakeholders that they have been breached or otherwise experienced a cyberincident?

What sanctions should policy-makers apply to compromised organizations?

11. Duty of assistance
How should public resources be drawn upon in the wake of a cyberincident?

12. Active defence
What technical measures should the private sector be empowered to use to deter and respond to cyberthreats?

13. Liability thresholds
What is the reasonable duty of care that an organization should have?

Who should bear the residual damages resulting from cyberincidents when an organization has sufficiently invested in security controls?

14. Cyberinsurance
What, if any, incentives should be offered to obtain insurance?

Which entities should be prioritized for these incentives?

Across these topics, a number of linkages and interdependencies exist. For example, an effective intelligence-sharing policy will help limit the spread of malicious software, and the greater adoption of encryption may limit the ability to monitor and police network traffic. In practice, what this means for business leaders and policy-makers is that cybersecurity policy-making efforts should be more collaborative and deliberative. Efforts should also be framed in the context of an ongoing iterative process rather than ad hoc and crisis-driven, resulting in patchwork legislation. Five key themes arise across the 14 policy topics covered by this document.

First, the acceptable scope of action for the public and private sectors should be more clearly defined. One manifestation of this issue is the question of where “safe harbour” provisions should or should not exist. For example:

- Policy around data and intelligence-sharing has been hindered by the absence of clear guidance for what constitutes protected industry collaboration.
- In the public-private context, the private sector has often been reluctant to share data with the public sector owing to concerns that revealed data will serve as the basis for future regulatory actions.

Second, the scope of permissible activity granted to security practitioners in the public and private sectors is often legally ambiguous at best. One common example of this difficulty arises in the context of cybersecurity research. In many jurisdictions, legitimate cybersecurity researchers — often colloquially called “white hat” ethical hackers in contrast to “black hat” malicious hackers — are uncertain as to the techniques and tools they are legally empowered to use to test systems. Furthermore, it is unclear how those researchers should inform others about security vulnerabilities. In one notable instance this past year in Hungary, in part owing to the absence of a legal framework around ethical hacking, an 18-year old was arrested after informing the Budapest Transit Authority about a vulnerability allowing customers to purchase online tickets at any desired price.⁴

Third, since digital traffic crosses national borders, a nation’s policy choices will usually have considerable impact on, and be impacted by, the choices of other nations. To help predict the longer-term effect of a policy position, it is worthwhile to consider the impact of a symmetric international policy response.

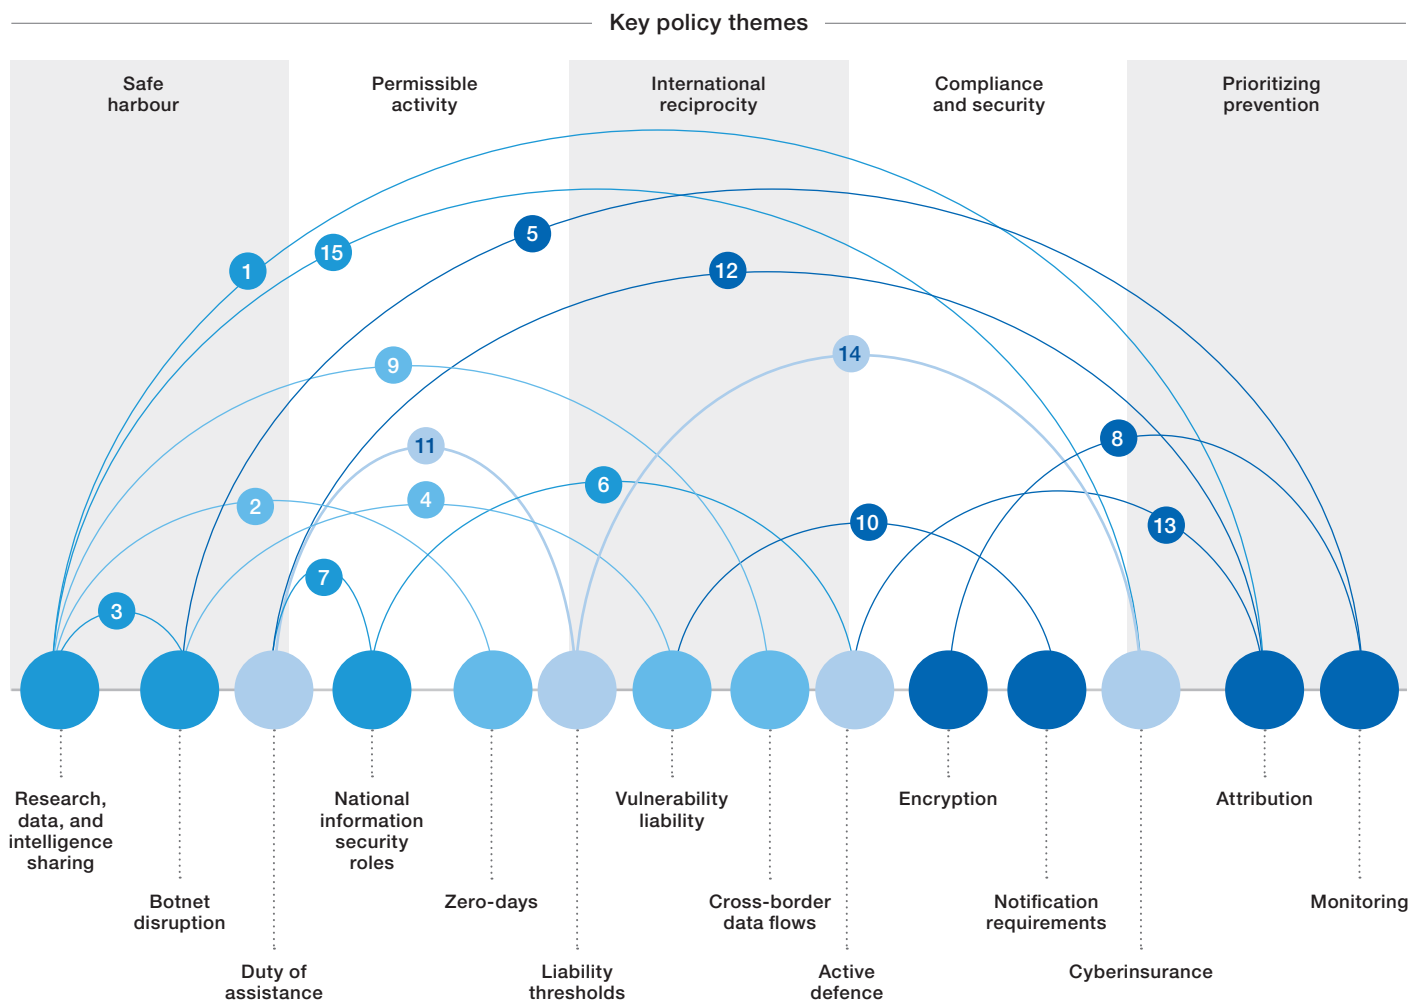
Fourth, in an effort to develop cybersecurity governance structures, policy-makers and, in particular, regulators, have begun exhaustively specifying processes and technologies for organizations to implement. Consequently, many organizations are devoting greater resources to achieve compliance. However, compliance may not necessarily advance cyber resilience. As more governments begin to formalize cyber-regulations, the costs undertaken by organizations to achieve compliance appear poised to grow.

Finally, for some policy questions, devoting incremental energy to developing preventive measures would avoid or limit more contentious trade-offs. For example, significant debate and intellectual energy has been devoted to discussing how software vulnerabilities should be disclosed. Considerably less policy guidance has been created to improve software coding quality standards. More secure software would reduce the stakes of the debate.

When considering the policy topic areas below, these cross-cutting issues should be taken into account while discussing each discrete policy option.

Cybersecurity policy landscape highly interdependent

Across fourteen major security topics, five key themes



Key linkages between policy topics

- | | | |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> 1 Attribution key element of intelligence, particularly for public sector 2 Zero-day vulnerabilities crucial opportunity for governments to share threat intelligence 3 Botnet disruption facilitated by rapid and well-coordinated research and action 4 Securing vulnerabilities through avoidance or patching may diminish threat surface for botnet operators 5 More invasive monitoring capabilities may allow ISPs to police botnet more effectively 6 Extent of active defence permitted by private sector key element of national roles and responsibilities | <ul style="list-style-type: none"> 7 Granular understanding of government duty of assistance fundamental to national cyber resilience 8 Greater adoption of strong encryption will hinder the ability to monitor network traffic 9 Limitations on cross-border data flows may introduce friction into intelligence sharing 10 Heightened notification requirements may result in increasing investment to secure known vulnerabilities 11 Duty to assist integrally linked with liability—where private sector cannot be reasonably expected to secure, government steps in | <ul style="list-style-type: none"> 12 Nation-state attribution may trigger government duty to assist the private sector 13 Active defence may result in collateral damage without well-defined attribution and safeguards (e.g. organization vs. nation-state) 14 Liability thresholds circumscribe the nature of cyberinsurance incentivized 15 Cyberinsurance can be more effectively priced and deployed given greater data and intelligence |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Note: List of connections between topics not exhaustive.

4.1 Zero-days

Definition

A zero-day vulnerability refers to an exploitable weakness in software that is usually unknown to a vendor. Since this vulnerability has never been shared publicly, no days have gone by to address the issue; thus it is on “day zero”. While such a vulnerability may be intentionally introduced, more often these vulnerabilities are inadvertently created. Zero-day exploits use an exploitable weakness to carry out an attack. Such exploits can be obtained through research and investigation or purchased through private-sector providers⁵

Policy model

Governments and private actors have increasingly debated the use of zero-day software vulnerabilities given their potential to militarize cyberspace.

To craft policy to address zero-days, it is necessary to understand that software vulnerabilities have a life cycle with three fundamental parts:

1. First, a zero-day vulnerability is written or introduced; a programmer writes code that is fundamentally exploitable or amenable to exploitation (e.g. software amenable to a vulnerable configuration).
2. Second, that zero-day vulnerability is discovered, typically after that software is implemented. Sometimes, different groups of researchers independently find the same vulnerability, often referred to as “rediscovery”.⁶ Researchers then develop means to exploit that vulnerability. These means, often called “exploits”, can then be purchased by either the public or private sector. The process by which the government decides to withhold or disclose vulnerabilities to a software vendor (with the expectation that the vendor develops a mitigation measure like a patch) is commonly referred to as the “vulnerabilities equities process” in the US.
3. Third, an exploit must be used on systems running the vulnerable software. To the extent that the exploit is deployed before mitigation measures (e.g. patches) can be developed, that exploit will result in some damage or harm. While patches are the most common form of mitigation measure, sometimes vulnerabilities are managed by other means (e.g. by avoiding a particular software configuration).

To address the first part of the life cycle, policy-makers may consider promoting or legislatively adopting coding standards for software vendors to limit the number of vulnerabilities created. Several commendable industry initiatives have been designed to promote secure coding standards based on expert guidance and community consensus (e.g. OWASP, NIST 800-64).^{7,8} However, even with greater resources devoted to coding standards and practices, zero-day vulnerabilities will continue to exist due to human error and other factors.

To address the second part of the life cycle, policy positions vary on two axes: involvement in the zero-day market and disclosure. (The third part of the life cycle is addressed under point 4.2 “Vulnerability liability”):


- Governments can completely exit the zero-day market and avoid research dedicated to finding software vulnerabilities. Alternatively, governments may choose to invest heavily in finding and exploiting vulnerabilities.
- Governments may choose to stockpile vulnerabilities for use at some future date or disclose these vulnerabilities to software vendors so that software vendors issue patches. The longer a given vulnerability exists, the more likely that it is rediscovered and exploited by other actors, including criminals and nation-state adversaries.

The risks and benefits of zero-day policy also depend on these two axes:

- The more heavily involved a government is in the research and purchasing of zero-day exploits, the more likely that software vulnerabilities are to be discovered generally. Government purchases incentivize researchers to find vulnerabilities, particularly as the monetary value of vulnerabilities rises. While there is an active community of researchers who do not have pecuniary motives for surfacing vulnerabilities for general research, private-sector vendors have begun to embrace the value of creating “bug-bounty” programmes where compensation is provided to individuals who inform companies about vulnerabilities in their products. As the public sector becomes a more active purchaser for knowledge of vulnerabilities — absent increased bounties — these programmes are less likely to be effective (as individuals will go to the highest bidder).⁹

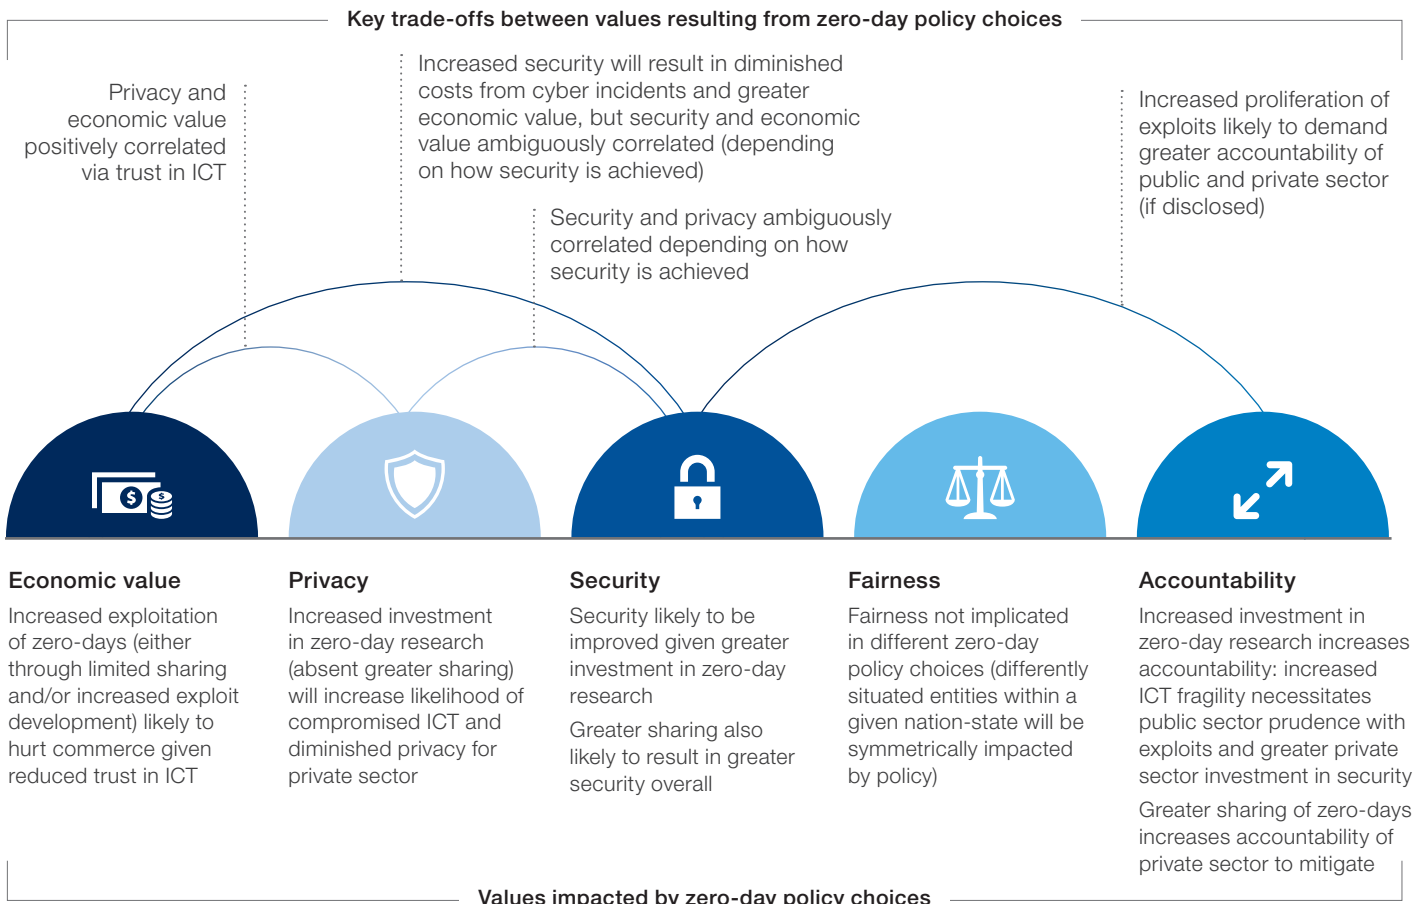
4.1 Zero-days

Policy model: Zero-days



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by zero-day policy



- The more government discloses vulnerabilities to the private sector, provided the private sector expeditiously mitigates those vulnerabilities, the more likely software is to be secure against all actors. This will have a double-edged effect as government may seek to utilize those exploits for law enforcement or espionage and will be unable to do so.

One important ancillary factor when considering a government's engagement with zero-day vulnerabilities is the placement of decision-making authority. In the US context, the "vulnerabilities equities process" is managed through the Executive branch owing to national security considerations. However, in other contexts, that same process could be an explicit legislative function. For example, continuing the US policy analogy, the American Congress could also pass legislation enumerating how vulnerabilities will be shared with the private sector.

Case study: Google, Project Zero

Google has embarked on an ambitious effort to find and disclose vulnerabilities to software vendors. Google's effort is noteworthy both in terms of its intellectual underpinnings and clever use of disclosure to incentivize patching vulnerabilities.¹⁰

Google has invested in discovering and disclosing zero-day vulnerabilities owing to a belief that patching zero-days is a highly efficient way to guarantee security. To put this belief in context, it is important to note that software is developed within the security methodology of defence-in-depth. A helpful physical analogue might be that of a medieval castle, which had many layers of defence, from moats to outer walls to ramparts and inner walls. The most devastating exploits utilize multiple zero-days to penetrate many layers of defence to breach a target (e.g. Stuxnet). Google believes that patching a small number of vulnerabilities can extraordinarily improve security as the impenetrability of one layer rebuffs an entire exploit.¹¹

Google has also pioneered an effective method to ensure that the vulnerabilities it discovers are promptly patched by software vendors. Namely, Google promises to disclose vulnerabilities to the public after a finite period of time to incentivize vendors to invest in developing and deploying patches rapidly.

Connecting policy to values

Zero-day policy choices balance a number of important trade-offs, particularly with respect to security, economic value, privacy and accountability:

- Security is improved in two almost symmetrically different scenarios. One line of thinking, more associated with the "offence" approaches in the framework, is that government can provide greater security for their citizens and firms by virtue of advance notice of an attack. Here, the government achieves advance notice of an attack by utilizing zero-day exploits against adversaries to predict or to thwart their intentions.
- Another way government can promote security is by sharing vulnerabilities with the private sector to "harden" information and communication technology (ICT), and to prevent adversaries from using those same exploits against a given country's citizens. The efficacy of this latter approach is premised on the timely private-sector development of mitigation strategies for shared vulnerabilities.
- The economic value associated with different zero-day policy scenarios is a function of a few different effects. First, greater security through both offence and defence is associated with less intra-state damages arising from cyberincidents. However, in the case of an "offensive" approach, security must be weighed against the risk of the same vulnerabilities being used against a country's citizens, as well (an instance where a zero-day vulnerability is "rediscovered" by an adversary). Additionally, some observers have noted that actions whose impact is to deteriorate trust in ICT create substantial intangible costs in terms of diminished ICT adoption.
- Privacy is also impacted by choices in some scenarios. Namely, in an "offence" approach, the improvement in security is premised on decreased privacy. In most cases, presumably decreased privacy should be limited to suspected criminals but the risk is that the confidentiality of innocent individuals will also be compromised.
- The extent to which vulnerabilities are shared with the private sector impacts the accountability of both the public and private sectors. If the public sector shares more vulnerabilities with the private sector, it is incumbent on the private sector to rapidly develop mitigation measures against those vulnerabilities, increasing the private sector's accountability. The opposite is also true; with more zero-day vulnerabilities held for greater periods of time, the public sector has greater accountability to ensure that those vulnerabilities are not weaponized by adversaries.

4.2 Vulnerability liability

Definitions

Known vulnerability — in contrast to a zero-day software vulnerability, a known vulnerability has been announced to the security community, generally with publicly documented methods to prevent its exploitation (e.g. patches or simple avoidance)

End-of-life — a product no longer supported by its developer with ongoing patches and updates

Open source software — software with public access to the source code of the program itself with licensing requirements. The source code is basically a list of commands that dictates how the program executes. Linux is an example of open-source software

Closed source software — software with proprietary and limited access to the source code of the program. Microsoft Office is an example of closed-source software

Software-as-a-service — a software distribution model in which one party hosts and maintains applications and makes them available to users over the internet

Policy model

Historically, the terms of use for licensing software have included some version of caveat emptor: “buyer beware”. Software vendors have explicitly avoided accepting liability for the damages caused by vulnerability exploitation. As software has become embedded more deeply into processes integral to an individual, business or even a nation-state, the potential damages associated with exploiting software vulnerabilities have also grown.

Moreover, there are tremendous swaths of legacy end-of-life (EoL) software for which vendors have entirely stepped away as they have moved on to develop newer, better versions.^{18,19} In some cases, the vendors no longer even exist. It is also increasingly the case that, in some software categories, market incentives and user choice cannot independently promote greater security given increased market concentration.

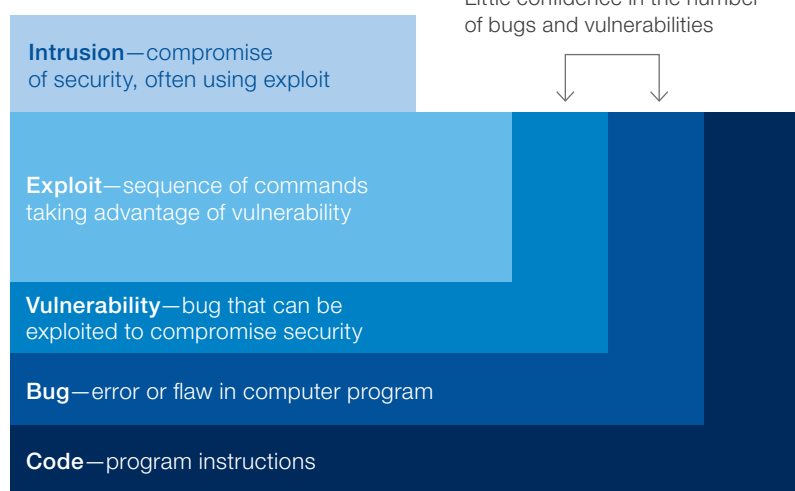
Liability can be attached to actors throughout the software ecosystem to calibrate incentives for security. One way to conceptualize how liability could be distributed is based on risk, where liability is determined on the basis of the potential consequences of exploiting a vulnerability (e.g. greater risk is associated with more stringent liability for a vendor). In thinking about assigning responsibility for securing vulnerabilities, two main questions emerge:

Who is liable or otherwise responsible for securing software? At least four broad sets of liability regimes exist, ranging from no liability to holding vendors liable for their software:

1. **No liability, code close-sourced** — this is the current norm where counterparties to software vendors may negotiate some level of accountability by exception. In this regime, if damages arise as a consequence of a vulnerability being exploited, the vendor is not held responsible.
2. **No liability, code open-sourced** — in exchange for being released of liability, vendors could be required to open-source underlying code. In theory, users and implementers of software would be more empowered to address vulnerabilities on their own. In this regime, if a vendor has open-sourced the code, the vendor is not responsible for consequences of software being exploited.
3. **User, implementer liable** — users and implementers could be held liable for damages arising from software being exploited. In practice, such a regime would create heightened incentives for users to contract for secure software. Within this context, the possibility also exists of differential liability that differs between enterprises with dedicated security teams and consumers (with more responsibility being attached to entities “that should know better”).
4. **Vendor liable** — vendors could be held liable for damages arising from software being exploited. For example, if a vendor did not issue a patch for a known software vulnerability, the vendor would be held liable if damages arose as a consequence (thereby heightening incentives for vendors to design and maintain secure software).


Breaches are the visible consequence of a very small share of the code base being exploited^{...2}

Very good sense of # of successful intrusions (i.e. breaches)



Good sense of code base

Size of codebase for popular software



Millions of lines of code

How should that liability shift when software transitions to EoL, or vendors go out of business?

- Given that most software vendors cannot afford to support software *ad infinitum*, there are justifiable concerns around attaching liability to software for perpetuity. To address those concerns, some commentators have proposed a sliding scale of liability, such that liability shifts as software enters EoL. For example, software may begin as a vendor's responsibility but as it becomes EoL, liability may be transferred to users/implementers. In such a regime, vendors would be held responsible for designing and maintaining secure software and there would be commercial incentives for users and implementers to upgrade to newer versions of software, when available.


Within this simplified framework, a number of hybrid arrangements could be proposed for splitting liability. For example, vendors can be held liable for rapidly providing mitigation guidance or a patch for a known vulnerability while users and implementers can be held liable for timely patch deployment.

Significant trade-offs are associated with different liability regimes:

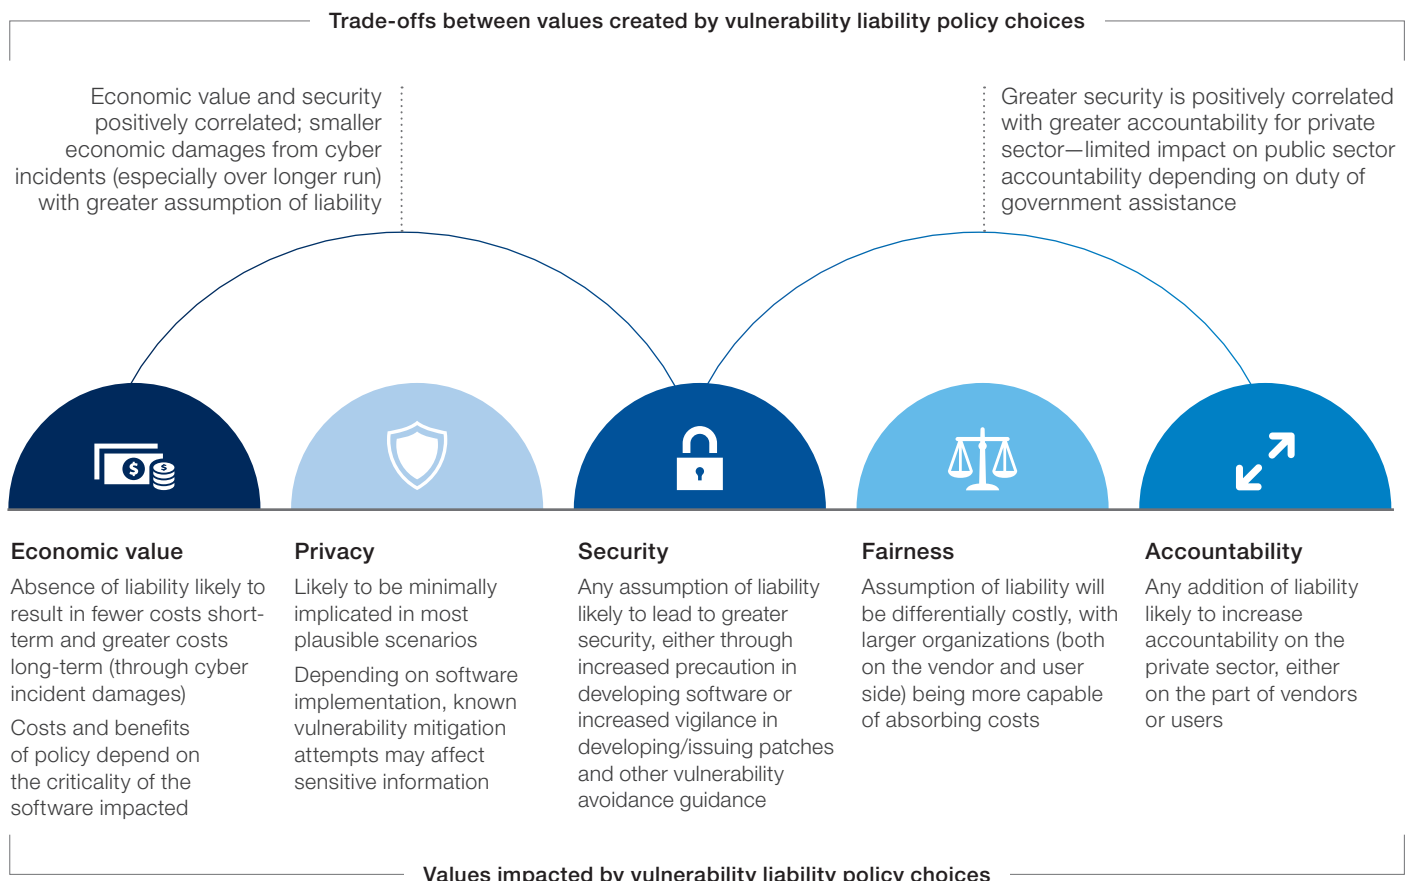
- **No liability, code close-sourced** — this is likely to result in fast software releases with significant security risks in the current environment.
- **No liability, code open-sourced** — this is likely to result in fast software releases with perhaps slightly diminished security risks. Put differently, it is not necessarily the case that open-source software is more secure. For example, in 2014, independent researchers discovered a cryptographic flaw in an open-source common implementation of encryption affecting two-thirds of the world's servers known as "Heartbleed".²⁰ Furthermore, open-sourcing code may dilute commercial incentives to innovate in software as a competing vendor may be able to engineer product changes more quickly relying on the investment of a first-mover.
- **User, implementer liable** — this is likely to result in slower software releases as commercial incentives from some users push vendors to develop more secure software.

4.2 Vulnerability liability

Policy model: Vulnerability liability



Key values trade-offs created by vulnerability liability policy



- **Vendor liable** — this is likely to result in even slower software releases as vendors will have every commercial incentive to provide security by design and deploy engineering resources behind maintaining secure software.
- **Embedded software** — bespoke software embedded within hardware that is not traditionally understood as a locus of computation (e.g. industrial control systems) — raises particularly difficult policy considerations. The depreciation horizon of the hardware often exceeds that of the software. Consequently, organizations with embedded software run systems that are often no longer supported or well-documented to realize value

The increasing adoption of software-as-a-service is addressing the issue of both EoL software and incentives to keep software secure. Vendors regularly update and provision software for customers (and generally maintain only a few versions).

Case study:

U.S. National Vulnerability Database, China National Vulnerability Database

To provide a structured repository for companies and researchers to mitigate and respond to security vulnerabilities, some countries have established national vulnerability databases. These vulnerability databases are then the source-of-record for vulnerability mitigation and form the basis of automated systems that security teams within enterprises use to prioritize patching.

Of course, a database of vulnerabilities also presents an opportunity for adversaries to develop exploits that weaponize vulnerabilities — at least until companies and researchers develop patches.

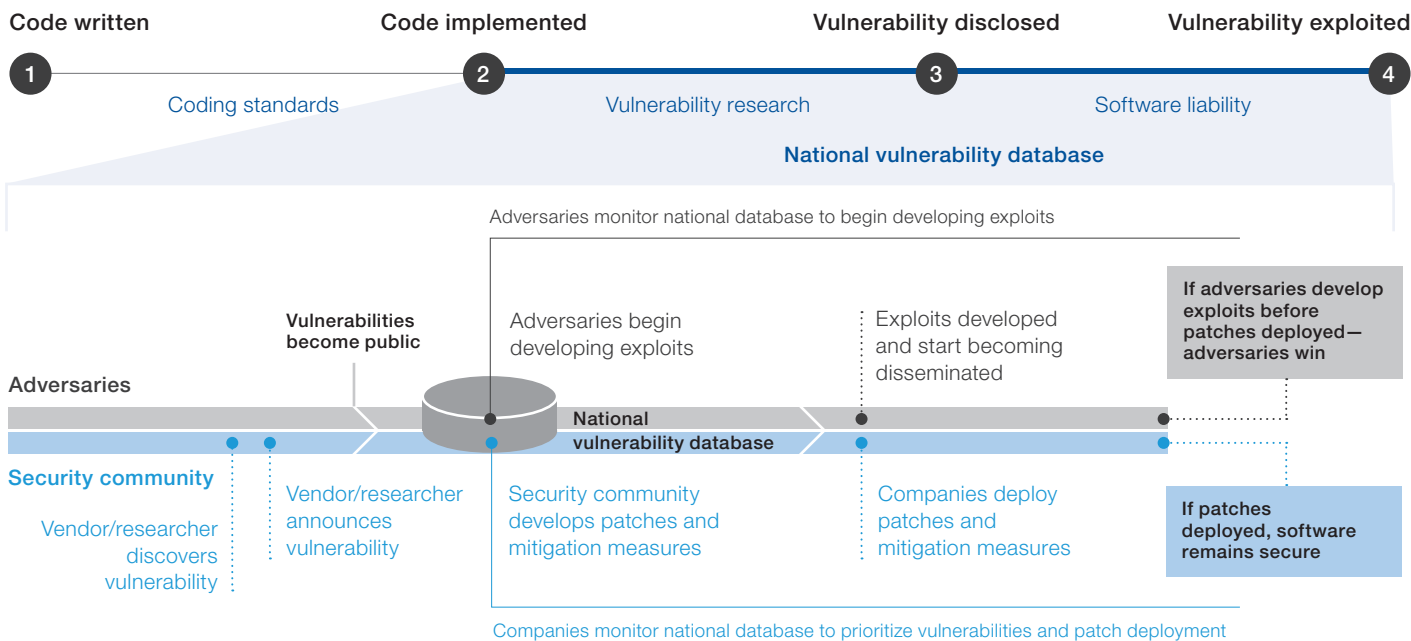
The United States has established a national vulnerability database, otherwise known as NVD, which serves as the international database of record with its nomenclature and structure. This database is constructed based on the voluntary submissions of vendors and researchers — a “push”. Often, this voluntary submission occurs some period of time after the vendor or researcher publicly discloses a vulnerability.

In contrast, China’s national vulnerability database, otherwise known as CNNVD, relies principally on a “pull” model and its researchers actively search for vulnerabilities surfaced by researchers, vendors and other sources, and document those vulnerabilities in the CNNVD. As a consequence, for the very same vulnerability, the CNNVD is often more timely than NVD. Recent research suggests that the average delay between first disclosure and availability on CNNVD is 13 days while on NVD the average delay is 33 days.

The practical impact of the staggered release of vulnerability disclosures of national vulnerability databases is the opportunity for a form of vulnerability arbitrage — adversaries learning about vulnerabilities on CNNVD and developing exploits before companies, particularly in the US context, have the opportunity to begin researching the mitigation of those exploits.

4.2 Vulnerability liability

Vulnerability disclosure



Connecting policy to values

Policy-making around attaching liability to those who build, implement and use software implicates three key values: security, economic value and accountability:

- Attaching greater liability around software use will increase its security. Whether liability rests with vendors or users, any attachment of liability is likely to increase security as the liable party undertakes greater precautions. If liability is associated with security vendors, vendors will develop more secure products. Indeed, such a requirement might encourage greater research into patching mechanisms that are both less intrusive and more difficult for users to avoid. On the other hand, increased liability for users will increase security not only through the user taking greater precautions (e.g. through more rigorously limiting access and control of particularly critical systems) but also by creating greater market incentives for vendors whose products meet more exacting security standards.
- Requirements to open-source software may have an ambiguous impact on security. While in theory the ability to have a community curate a codebase and test it would improve security, there is little empirical evidence that open-source

software is inherently more secure than its closed-source equivalent. That said, requirements to open source unsupported software that was previously closed source may improve security. For software vendors selling closed-source software, such a requirement would prevent those companies from monetizing products while avoiding the responsibility of ensuring security (open-source software is generally monetized differently from closed-source software). As a consequence, vendors may be more greatly incentivized to provision and secure products for a longer period of time.

- Insofar as greater liability results in greater security, the economic value of greater liability is positive. However, there is not a simple linear relationship between greater liability and greater economic value; at some threshold, greater liability will impose significant costs on the software ecosystem, outweighing the mitigated security damages. Furthermore, assigning liability may decrease the speed of innovation as vendors now bear the equivalent of a “warranty cost”, either directly or indirectly, through the demands of their users.
- Greater liability is likely to lead to greater private-sector accountability, particularly if vendors are held liable for security directly.

4.3 Attribution

Definition

Attribution — determining the identity or location of an attacker or an attacker's intermediary. In the case of cybersecurity, attribution is a particularly difficult problem as adversaries can mask their identity or even originate attacks from deceptive and unwitting locations (e.g. using a hospital's network as a staging ground)²³

Policy model

As cyberspace has become increasingly weaponized, determining the perpetrator of an attack to impose costs on the attacker and prevent future attacks has become more important. In contrast to traditional crime, in many contexts, this determination is the result of private actors responding to a cyberincident, which is particularly salient when private actors accuse nation-states of criminal activity.

A key policy question on attribution is: how should government engage with the private sector when the private sector publicly alleges that a particular actor is responsible for a given attack? In private, for purposes of research and intelligence gathering, attribution — connecting an alleged adversary to a given attack — has limited potential consequences. Furthermore, attribution is core to the functioning of researchers and security teams: knowing that a particular adversary is likely responsible for an intrusion enables drawing upon documentation on the historical tools and techniques used by that adversary to respond more quickly to an incident.²⁴

Policy stances on attribution principally hinge on two positions: the government's obligation to respond to a claim of attribution and the government's validation of a particular company's attribution of an attack to a particular adversary:

- Governments can have a standing policy where no obligation arises out of attribution. In practice, this would mean that if a company asserted that a given actor, whether a state or an individual, attacked an entity, the government would have no affirmative obligation to act on that assertion. Alternatively, government could be obligated to respond, and at least investigate credible claims of an attack against one of its citizens by a foreign or domestic actor.
- When the private sector makes public claims about the identity of a given attacker, governments have two choices: to affirm and (in)validate a claim or to avoid public comment.

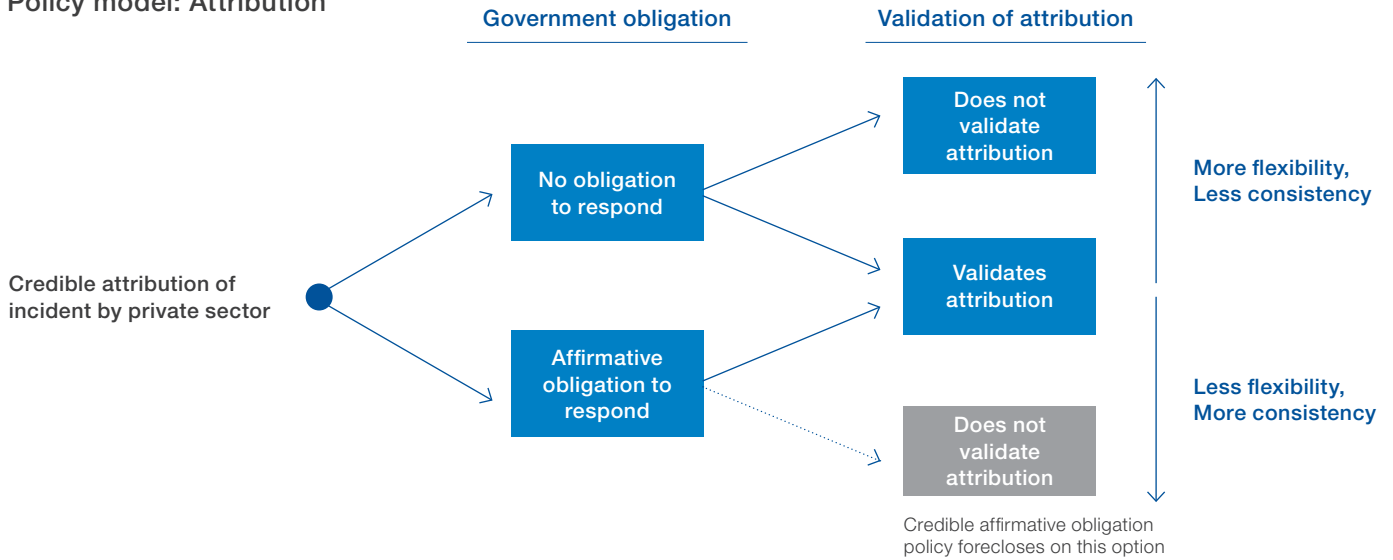
The risks and benefits of policy also vary on these two axes:

- If a government's policy is that no obligation arises out of attribution, then there are limited short-term potential collateral consequences if a company asserts that a particular actor is responsible for a given crime, with less opportunity for an incident to escalate into a diplomatic issue. In the long run, however, failure to attribute an attack could undermine a country's deterrence posture, thereby inviting future attacks and undermining public confidence. Additionally, in the absence of government reaction to attribution, efforts to coordinate research on the actors behind a given attack may be delayed. Where a government has an affirmative obligation to act on attribution claims, the potential short-term collateral consequences are magnified. For example, if a state is accused of perpetrating an attack, the host state may risk worsening diplomatic and economic relations with the alleged attacker state if it affirms the attribution. The host state may also reveal capabilities or vulnerabilities that are better kept concealed. In the long run, however, attribution may improve a country's deterrence posture, thereby limiting future attacks and building public confidence.
- A policy of validating private-sector claims of attribution risks private companies being effectively considered as government appendages, hampering the capacity of some businesses to operate outside of a given country (given associations with a national government). Furthermore, such a policy is fundamentally impracticable in the long run for multinational organizations. In the hypothetical case of country-related claims of attribution, if a company operates in 100 countries, any single country's insistence to validate claims of attribution could be imperilled by a reciprocal differing response abroad. Multinationals are then forced to pick between customers and national demands.

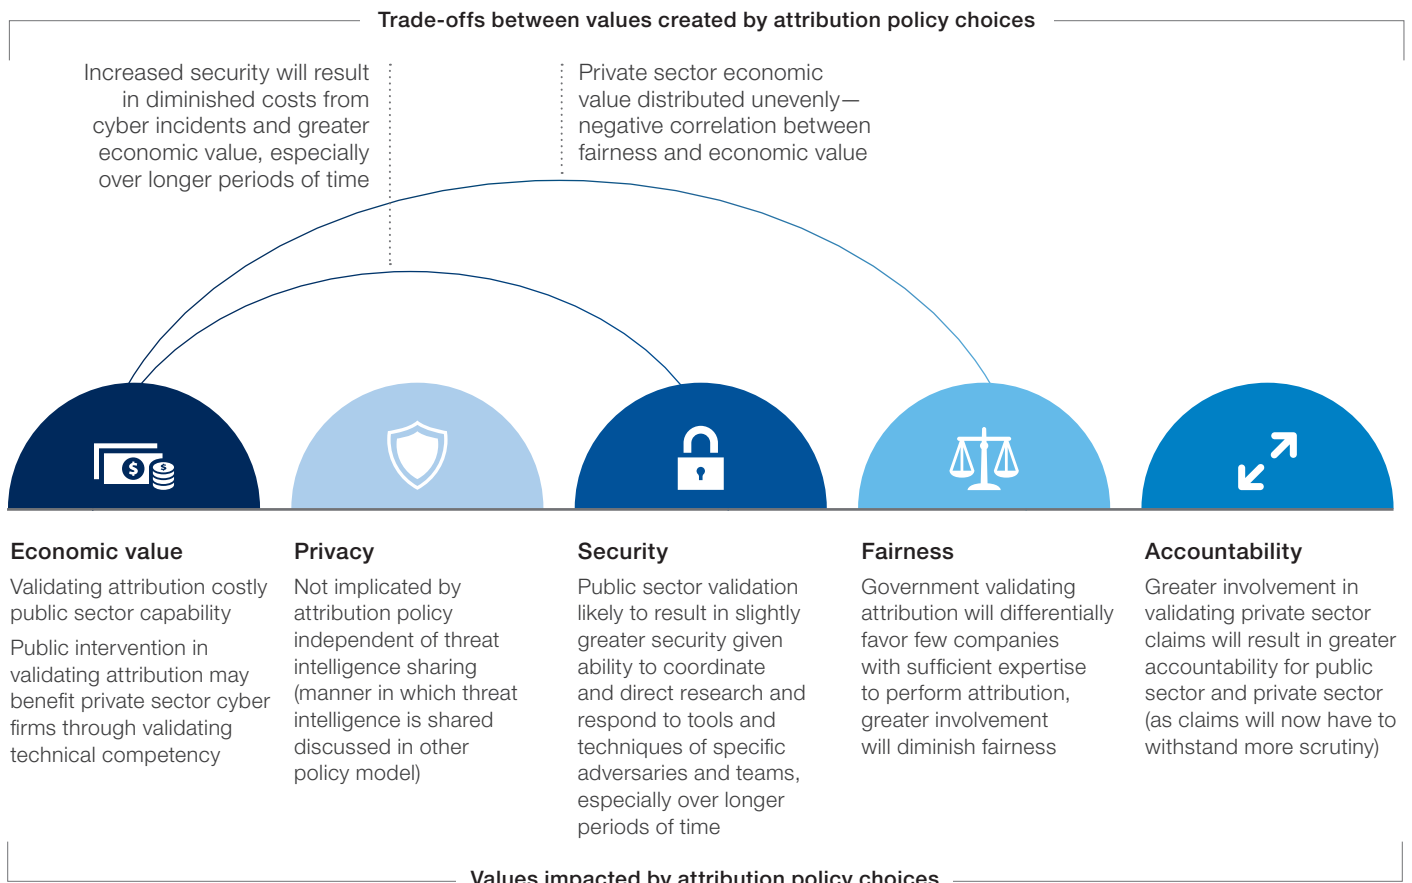
Most commentators agree that while attribution is technically possible, in practice few private-sector actors have the capabilities to reliably establish it, and many are headquartered in the United States. The reliance on private-sector actors to engage in attribution, particularly given the geopolitical risks, may result in a system brittle to accusations of nationalism clouding judgement.²⁵

4.3 Attribution

Policy model: Attribution



Key values trade-offs created by attribution policy





Connecting policy to values

Attribution policy brings into high relief certain trade-offs between security, economic value, accountability and fairness:

- Increased public-sector validation of private-sector attribution claims may improve security over the long run, depending on how such a policy is implemented. Greater private-sector firm awareness of how specific teams use particular tools and techniques to compromise networks will help inform efforts to develop technology and processes to mitigate these measures. However, of note is that the security improvement is to a greater extent contingent on understanding how specific adversaries operate rather than on the nation-state component of attribution itself, which is of limited practical value for most security practitioners.
- The economic value of public-sector validation of attribution claims is ambiguous in the short run and positive in the long run. In addition to reducing cyberincident costs, public-sector validation will financially reward the few private-sector firms capable of establishing attribution as a form of “approval” testifying to the accuracy of a given firm’s work. But costs are also associated with building sufficient and sustainable attribution capacity in government and, in some circumstances, public-private sector collaboration may impact perceptions of a company’s independence.
- An increased role for the government in responding to private-sector claims of attribution will increase accountability. The government’s heightened responsibility will not only increase its own accountability but also that of the private sector, whose attribution claims will be scrutinized. The private sector will either improve its own attribution capabilities, or it may defer entirely to the government to avoid both the costs and risks of being incorrect.
- However, an increased role for public-sector validation will decrease fairness both in terms of security and economic value. Very few security teams have the operational capabilities to practically benefit from the public sector investigating and sharing the tools and techniques used by adversaries, particularly nation-states. Additionally, very few firms are able to establish an adversary’s identity. Those firms may be differentially financially rewarded by the market for proof of their capabilities affirmed by the public sector.

4.4 Research, data and intelligence sharing

Definitions

Threat intelligence — insight into the capability and intent of an existing or emerging menace. In the context of cybersecurity, this could range from technical indicators (e.g. samples of suspected malware) to non-technical indicators (e.g. hacker forum discussions)

Personally identifiable information (PII) — any data that could potentially identify a specific individual; any information that can be used to distinguish one person from another and can be used to de-anonymize anonymous data can be considered PII. Breach notification laws typically focus on notifying the public when PII might have been exposed to unauthorized individuals, particularly in the context of financial or medical information

Policy model

As many more organizations in the private and public sectors are subject to cyberattacks, both sectors have been seeking to develop structured collaboration to ensure that individual research, data and threat intelligence are pooled to create a collective immune system-like response.

The key policy question regarding threat intelligence is: what is the government's role in sharing threat intelligence? Coordination and sharing are necessary as individual actors rarely see the entire landscape of potential threats. Threat intelligence has historically existed within a fragmented landscape, with companies relying on a combination of private-sector feeds provided by security vendors and internal research, with limited public-sector involvement.

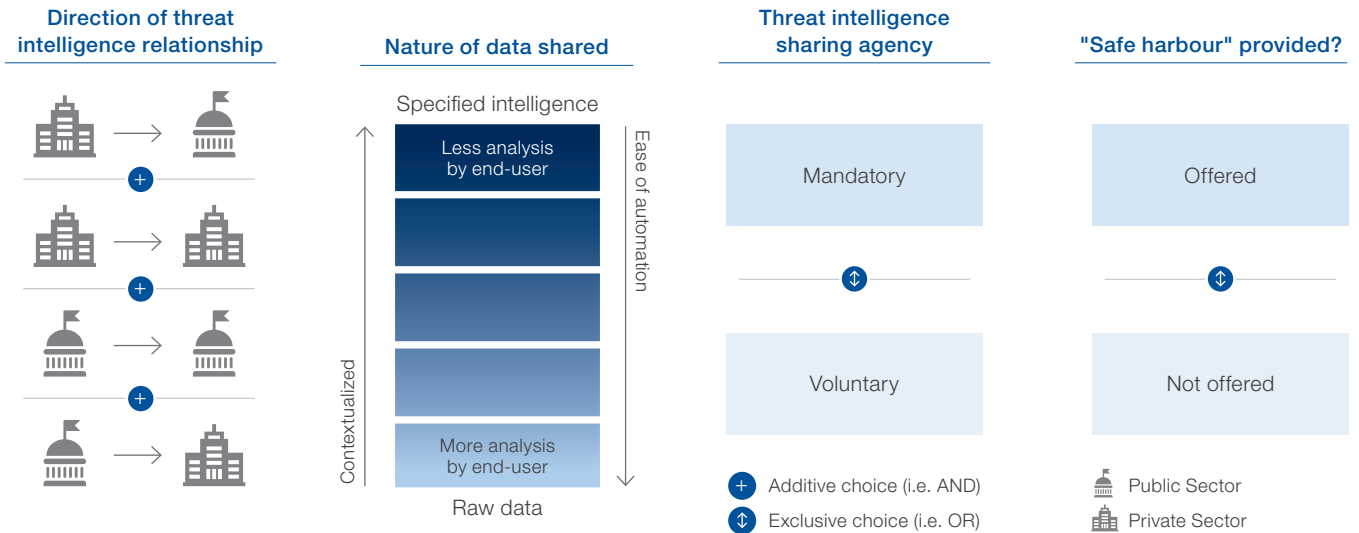
In the last few years, however, governments have increasingly attempted to formalize threat-intelligence-sharing relationships between the public and private sectors and to create scalable models for sharing data without compromising sources and methods (in the case of government-provided threat intelligence) or privacy (in the case of private-sector-provided intelligence).

Some currently proposed regulations intended to promote privacy and the limited sharing of PII may actually hinder information-sharing relationships. Companies may have legitimate concerns regarding whether collaboration will create more legal liability than averted cyberdamages.

Policy positions on threat intelligence must consider four major questions:

1. Who is involved in an intelligence-sharing relationship? Different models have been pursued and calls have been increasing to establish a broader direct-sharing relationship not only among the private sector but from the public to the private sector.
2. What will be shared? Everything from raw data (e.g. URLs) to analysed intelligence (e.g. URLs that are determined to originate suspicious traffic using a specific protocol to target a particular vulnerability with the aim of exfiltrating a particular type of data) can be shared. However, the more analysed intelligence becomes, the less automatable its sharing becomes. Automation pertaining to threat intelligence is important because cyberattacks operate at network speed — responding quickly and updating firewalls and malware signatures may be decisive in preventing an intrusion. To put this into context, a recent report measures the median “dwell time” of cyberattackers, the length of time an attacker is within a network, as 99 days.²⁶ While increased automation can diminish this time, the speed that automation can provide is not a panacea. Automating a response into one's security posture may impede the legitimate use of an application or access to particular data.
3. Is sharing mandated? Governments can choose to allow threat intelligence sharing on a voluntary or mandatory basis.
4. What safe harbour will be provided? The concern here is specific to the private sector; namely, companies would like to avoid incurring customer or regulatory liability for sharing threat intelligence. In the United States, for example, companies historically were concerned about claims of anticompetitive collusion whose basis would be threat intelligence sharing.²⁷ As such, a safe harbour from liability is often attached to a mandate to share intelligence.

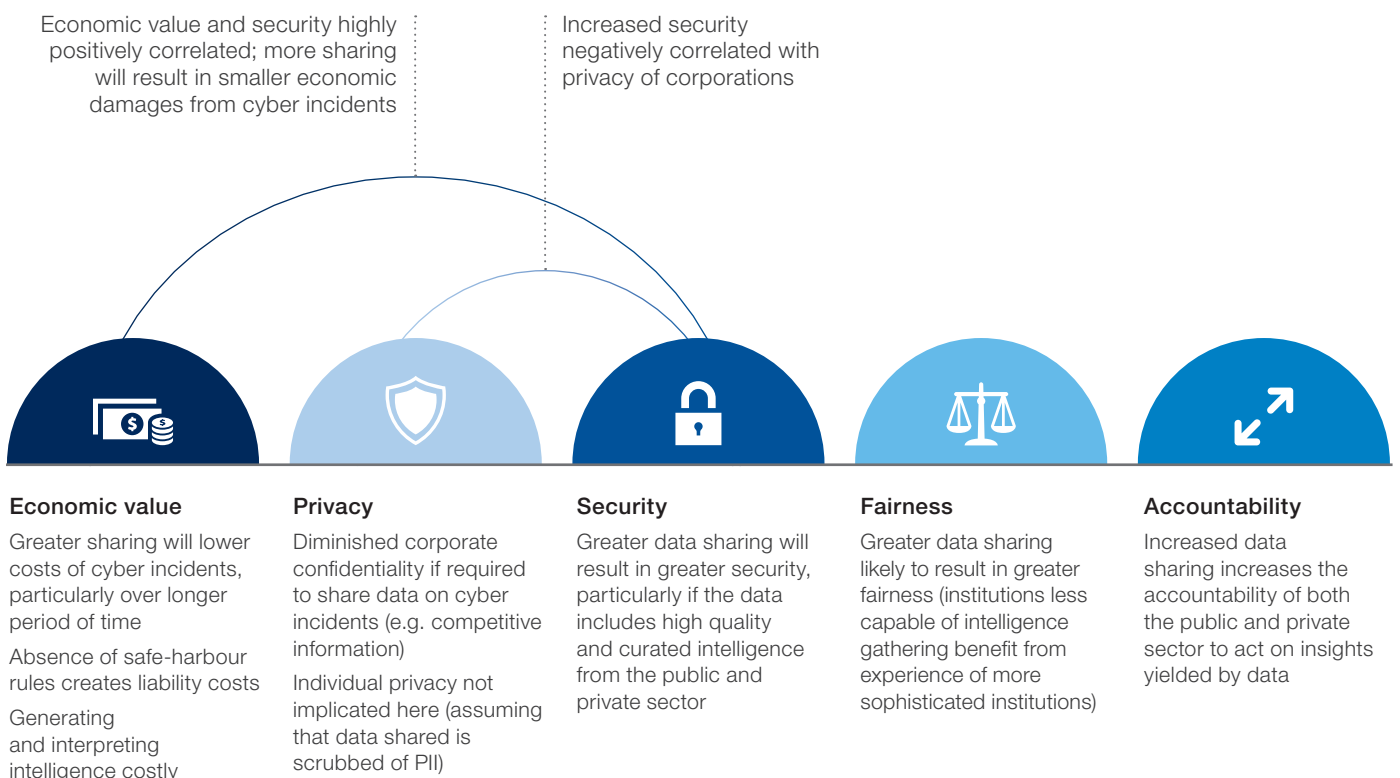
Policy model: Research, data and intelligence sharing



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by intelligence sharing policy

Trade-offs between values created by intelligence sharing policy choices



Values impacted by intelligence sharing policy choices

4.4 Research, data and intelligence sharing

The risks and benefits to the different arrangements for each of the aforementioned questions are significant:

- The greater the number of participants, the more threat intelligence can be generated, shared and validated.
 - While greater data volumes are not necessarily correlated with greater capability to generate insight, in theory more data provides a richer sample for companies to analyse and draw inferences from. In some circumstances, increased data that is not properly curated can impair a security posture as some participants may contribute less-actionable or lower-quality intelligence.
 - The richer the intelligence shared, the more actionable it is for practitioners to secure their own organization's systems against a particular threat. Of course, for those generating such intelligence, documenting an adversary's motivations and providing contextualized analysis that another company or the government can act upon requires significant resources.
- Mandates are likely to result in greater volumes of data being shared along with concomitant costs. In addition to the prior concerns about the diminishing and even negative returns to increasing volumes of data, mandated formalized sharing may reduce informal and productive arrangements developed by the security teams of larger, well-established institutions (particularly in the closely knit international financial sector).
 - National policy on threat intelligence sharing must be sensitive to international concerns and the implications of potential reciprocity. For example, compelling a multinational company to share threat intelligence with the public sector could imperil a company's international business if international customers feel that privacy or confidentiality may become compromised.



Case study: Department of Homeland Security, Automated Indicator Sharing (AIS)²⁸

To promote the rapid and timely sharing of threat intelligence indicators between the public and private sectors, the U.S. Department of Homeland Security (DHS) created a voluntary and automated cyberthreat-sharing programme to facilitate collaboration between the public and private sectors. The DHS programme is a remarkable innovation in the following key respects:

1. AIS facilitates sharing between the public and private sectors, addressing a common refrain from large companies that intelligence-sharing relationships often appear one-sided.
2. AIS is automated, such that threats at network speed can be addressed almost as quickly as they materialize.
3. AIS has vitiated the principal confidentiality and privacy concerns surrounding the use of automated threat intelligence sharing by providing limited safe harbour and scrubbing threat intelligence for PII.

One lesson from AIS, however, is the difficulty of obtaining traction for any voluntary threat intelligence sharing programme, at least when such a programme is “sub-scale.” Like any network-based model, the marginal value derived from AIS is small for the first few participants, even given the public sector’s contribution. However, as AIS becomes broadly adopted, each marginal would-be participant would likely derive greater value and thus more would join (a version of a “flywheel” effect).

Connecting policy to values

Research, data and intelligence-sharing policy implicate a number of values including, most importantly, security, economic value, accountability and fairness:

- Greater data sharing is likely to lead to greater security, particularly over the longer term. In the short run, greater data sharing may have an ambiguous impact as security practitioners learn to use and deploy analytical tools to ingest and process more data and draw insights. However, in the long run, as simpler forms of data analysis become automated and accessible tools augment human reasoning, greater data sharing should lead to greater collective security.
- Greater security will be realized over the longer term, meaning that the economic value of reduced costs from cyberincidents will also be realized over the longer term. That said, in the short term, greater data sharing is likely to impose significant capital and operating expenditures. Not only will sharing at machine speed require investing in new tools, but intelligently leveraging these tools and training new cyberprofessionals to use and embed them as part of the security workflow will be costly.
- Greater data sharing generally increases accountability for all ecosystem participants. Entities in the public and private sectors will need to take responsibility both for contributing actionable intelligence and for acting on the intelligence shared by ecosystem participants. Divergent sharing models, for example mandating the private sector to share intelligence with the public sector without a reciprocal demand, result in differential accountability. In the example, the private sector has increased accountability whereas the public sector does not.

4.5 Botnet disruption

Definitions

Botnet — a term derived from the words robot and network, a bot is a device infected by malware that becomes part of a network, or net, of infected devices controlled by a single attacker or attack group. The botnet malware typically looks for vulnerable devices across the internet, rather than targeting specific individuals, companies or industries. The objective of a botnet is to infect as many connected devices as possible, and to use the computing power and resources of those devices for automated tasks that generally remain hidden to the users of the devices²⁹

Botnet takedown — successfully taking permanent control of the entirety of a botnet or otherwise rendering the botnet useless

Botnet disruption — partially impairing the operations of a botnet to diminish its impact. In recent years, given distributed communication and organization methods, it has become more difficult to fully disable a botnet (a takedown)

Policy model

Botnets have been a persistent threat and problem confronting policy-makers as the internet's ubiquity has increased. In recent years, the spectre of this threat has grown symmetrically to the exponential growth in connected devices, known as the internet of things (IoT), and the internet traffic they generate. And given the tremendous promise of IoT, policy-makers are scrambling to structure policy that promotes IoT adoption without compromising security and trust.

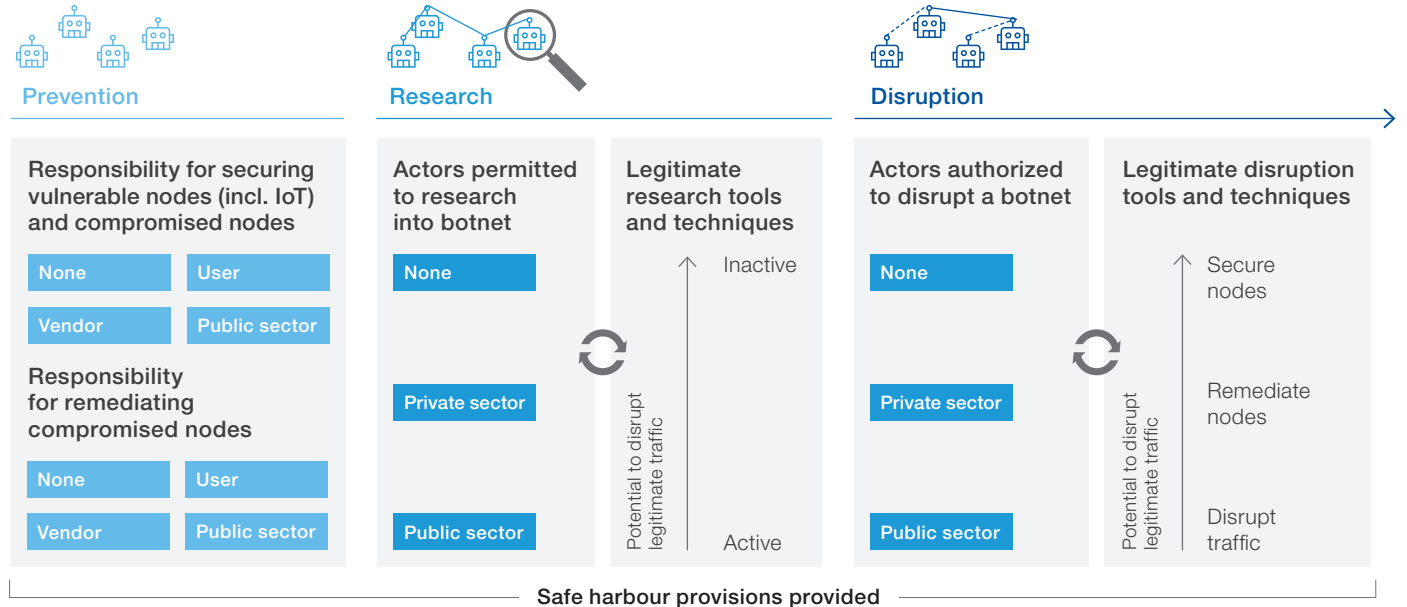
The key policy question related to botnets is: what degree and form of intervention is appropriate to prevent, research and disrupt botnets? In understanding how to manage the growing threat botnets pose, it is analytically helpful to divide policy into three questions across the life cycle of a botnet, from creation to disruption. At each stage of the policy discussion, the "safe harbours" provided for actors must be kept in mind. Close collaboration is necessary between the public and private sectors on this issue in particular, and good faith efforts may have unintended consequences:

1. What can be done to prevent the proliferation of botnets? The question of prevention involves two components. First, policy-makers must clarify responsibilities around remediating known, existing botnet nodes. Is there any responsibility attached

to using a connected device known to be part of a botnet, and is it the user's responsibility, the vendor's responsibility, or is the public sector responsible for ensuring that affected devices are patched? Second, the question also applies in terms of preventing the creation of new botnet nodes. Policy-makers may undertake to allow the private sector to police itself using market incentives or regulate minimum security standards (particularly for IoT).

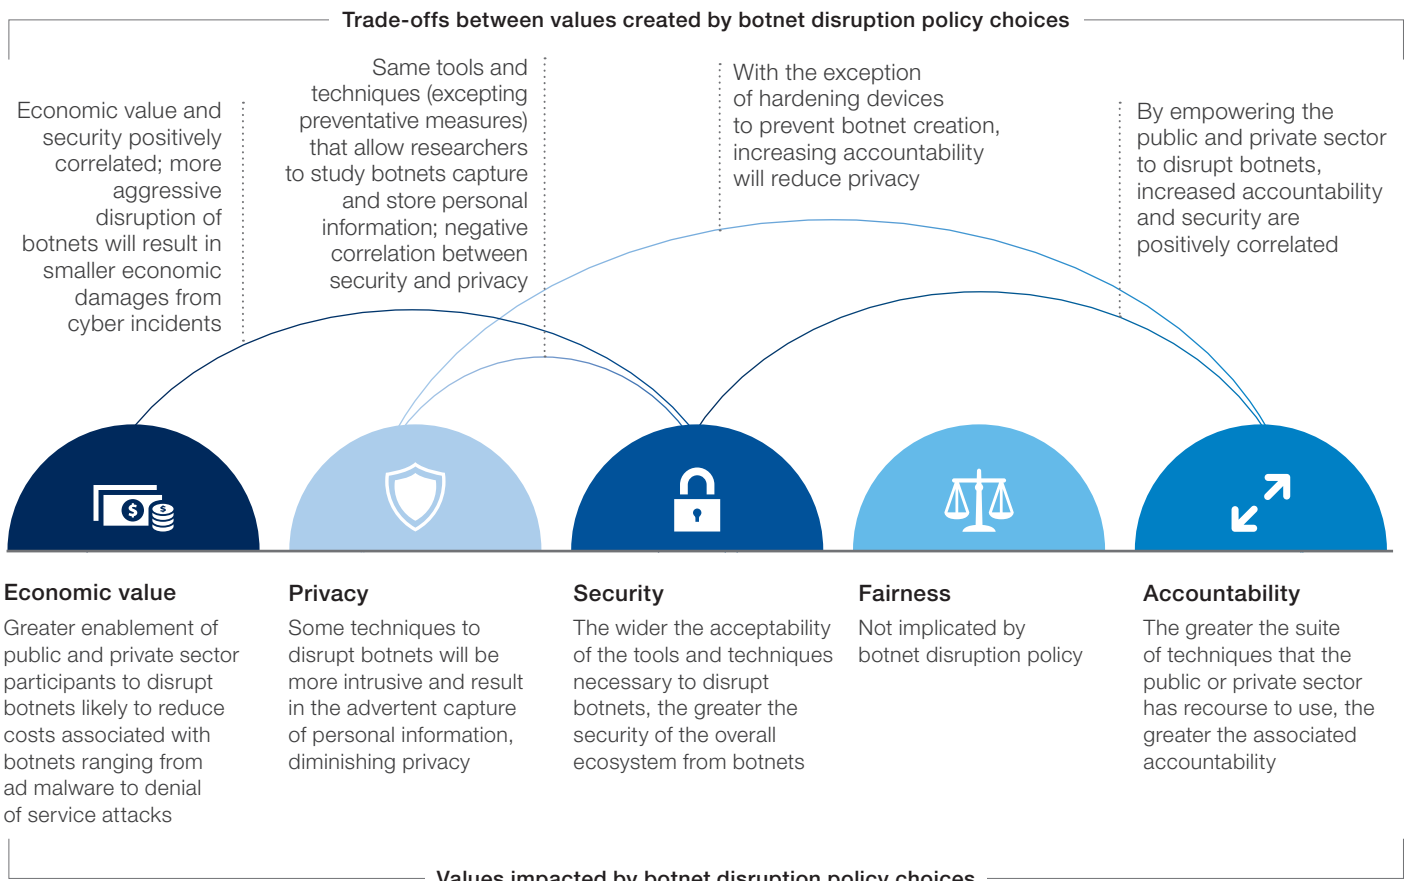
2. How can existing botnets be researched and studied? Who is allowed to research botnets and what techniques are they allowed to use? Methods to study botnets (to disrupt them) often require personally identifiable information and reshaping network traffic. For example, researchers may direct traffic from a known botnet node to a server for analysis. The traffic from that node, in addition to containing illegitimate or malicious traffic, may very well contain legitimate queries to websites containing sensitive information. As a consequence, it is important to clearly outline who is allowed to undertake these actions, especially as these methods may disrupt legitimate traffic.
3. How should actors throughout the ecosystem undertake disrupting botnets? Similar to research into botnets, tools and techniques to disrupt botnets have the potential to negatively impact legitimate users and their day-to-day well-being. As such, it is important to outline who is allowed to disrupt a botnet and what methods they are empowered to use. Certain methods have greater potential consequences associated with them than others. Techniques that attempt to remediate individual botnet nodes (e.g. removing malware from individual nodes) rather than disrupt the traffic between nodes are inherently less likely to create further network traffic issues. On occasion, techniques to disrupt traffic between nodes result in the disruption of legitimate traffic, as well.³⁰ Recent botnets have brigaded a number of IoT devices, including smart TVs and webcams. Attempts to disconnect these devices from the botnet may render them unusable. If the unusable IoT device is a smart refrigerator, disruption may simply be an inconvenience. In the future, if the botnet IoT device is an embedded sensor in an industrial control system, disruption may impact the power grid. As such, it will be increasingly important to undertake measures to mitigate collateral consequences from botnet disruption by understanding the nodes of a botnet more thoroughly while respecting the privacy concerns that may arise from the necessarily more complete perspective of node traffic.

Policy model: Botnet disruption



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by botnet disruption policy



4.5 Botnet disruption

The risks and benefits associated with policy positions on each of these questions are significant:

- Attaching no liability for securing compromised devices or would-be bots is likely to result, at least in the short term, in a proliferation of new devices with less security. However, in the medium to long term, given limited market incentives, many vendors will choose to architect security as an afterthought. So far, little market evidence supports the proposition that consumers will attach monetary value to secure IoT devices.³¹ Put differently, there is limited evidence of a “best-of-both-worlds” scenario where market incentives promote sufficient security.
- Placing liability in the hands of users is likely to result in similar outcomes as a situation of no liability. More savvy users (e.g. enterprises) and persistent consumers will resolve security issues but many will avoid or be unable to deploy security solutions.
- Increasingly, regulators are exploring a combination of demands on internet service providers (ISPs) and minimum standards for IoT devices to promote security. In contrast to an environment of no liability, this arrangement is likely to slow down the development of IoT devices and impose greater costs on ISPs, but will promote greater security.
- Another way to understand the risks associated with botnet policy, particularly as it pertains to IoT regulation, is from the standpoint of risk management of IoT device vendors. Vendors must balance the risks associated with business competition, regulatory overreach and reputational damage.

Case study: The internet of things, botnets and denial of service attacks

In recent years, the increasing proliferation of IoT devices has served to fuel ever-larger botnets whose network traffic can be redirected towards targets to overwhelm their ability to respond to network queries and denying legitimate users access to internet services. For example, the Mirai botnet in 2016 indirectly resulted in accessibility issues for major websites. To help craft policy to address this issue, it is helpful to understand why IoT devices are relatively vulnerable to being brigaded into botnets, and potential policy and technical solutions.³²

1. Why are IoT devices relatively vulnerable? IoT devices tend to have a few attributes that make them especially vulnerable to becoming part of a botnet:
 - **Diversity of devices** — the software ecosystem supporting IoT devices is more heterogeneous than PCs or smartphones, providing a greater exploitable threat surface.
 - **Limited computing resources** — IoT devices are often little more than sensors connected to the internet. As such, they lack the computing power to run conventional security protocols that often consume 30% of a typical laptop’s computing resources.
 - **Network persistence** — IoT devices are purpose-built to be able to connect to the internet in almost all circumstances. Thus they are ideal for bad actors to take control over the internet, as well.




- **Passive use case** — Unlike traditional computing resources (where abnormal processes deleteriously impact a user's experience), the mainly passive use of IoT devices means that, often, few outward signs or behaviours indicate malicious action is being taken with a given device's network traffic.
2. What are the technical and policy solutions promoting greater IoT security? Regarding the policy options to address botnets, distinguishing between different architectural approaches is helpful. In each case, the stakeholders and technical challenges are different:
- **Hardening individual nodes** — new technologies are emerging to secure devices with a “thin” agent (an agent that does not consume as many resources on the endpoint).
 - **Securing the connection between nodes and the internet** — several vendors have released what is known as a “proxy”, which acts as a barrier between the nodes and the internet, filtering and monitoring traffic for abnormalities.
 - **Monitoring internet traffic** — ISPs are increasingly using technologies to “scrub” network traffic patterns to look for telltale signs of botnets (e.g. coordinated network behaviour).

Connecting policy to values

Policy choices around preventing, researching and disrupting botnets create important trade-offs between a number of values, including economic value, privacy, security and accountability:

- Greater enablement of public- and private-sector entities (including academic researchers) in researching and disrupting botnets will likely improve security and subsequently reduce security-related damages. At the same time, greater measures to prevent botnet creation through attaching liability are likely to create short-term costs, both in terms of more expensive devices and fewer devices being adopted by users. On net, more “aggressive” policy around botnets is likely to create short-term costs that will be outweighed by the long-term benefits.
- Enabling more entities to use more invasive techniques to research and disrupt botnets will increase the private sector's accountability. These techniques, used incorrectly, threaten the well-being of innocent bystanders and as such are associated with the greater private-sector obligation to act responsibly. Furthermore, as the private sector is more empowered to act, the continuing presence of botnet nodes will be less acceptable.
- However, the more entities are empowered to act and the more intrusive the techniques they are allowed to use, the more privacy will be diminished. The very techniques that allow researchers to determine whether nodes are acting as part of a botnet involve capturing data coming to and from those nodes that may be sensitive.



4.6 Monitoring

Definitions

Metadata — basic information about data, which can make categorizing, finding and working with particular instances of data easier; in the case of surveillance — especially on the part of government agencies — metadata not only facilitates categorizing and retrieving content but provides information on its own and may also be used to legitimize collecting and examining content

Internet service providers (ISP) — companies that provide access to the internet and other related services, such as website building and virtual hosting³³

Technology platform — companies that facilitate communication or messaging over a variety of protocols (e.g. mobile messaging, instant messaging, email, etc.)

Policy model

One way to frame this policy question is similar to the discussion on encryption; namely, at a fundamental level, who should be able to see what? While end users necessarily observe digital content, what content should others be able to monitor to promote security and other valid national interests (e.g. privacy)?

This tension surfaces in at least three scenarios: between an employee and an employer, between a customer and an ISP, and between a user and a technology platform.

Additionally, for purposes of analytical simplicity, it is helpful to separate internet traffic into two components. The first is metadata, the instructions that allow entities to understand to whom to address content and how to relay it. Metadata is intrinsically difficult to mask (e.g. through encryption) — for example, if the address of the recipient of data is masked, how will an ISP know who to transit that data to? The second is content, the actual digital payload that a user is perceiving. Content may include malware and other malicious digital payloads.

In the context of monitoring metadata and content, a wide variety of policy options can be undertaken, but it is helpful to think about three choices: what is approved or a priori legitimate, permissible (e.g. by court order) or forbidden (never permitted). To take a few examples:

- Government may, by exception, be permitted to monitor metadata in the investigation of a crime (e.g. under subpoena). Alternatively, policy-makers may choose to entirely abrogate the government's ability to perceive any digital data in transit by limiting the gathering of metadata, in general.
- Employers may be presumptively allowed to observe the digital content accessed by employees, particularly if employees are utilizing employer-provided resources.

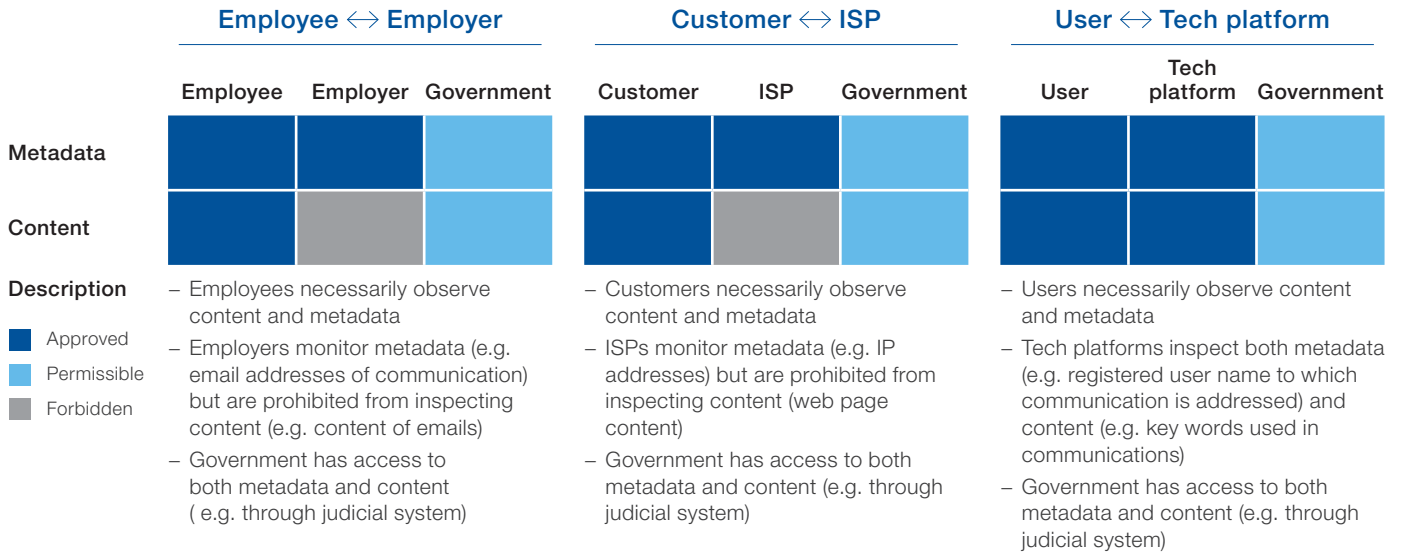
Each policy configuration has its own unique risks and benefits, but a few are common:

- The greater the extent to which monitoring content (that may include malicious payloads) is limited to users, the more end users are responsible for their own security; that is to say, security measures that could otherwise be deployed by government, a tech platform, an ISP or an employer cannot be utilized.
- Placing increasing capabilities and responsibilities to monitor content in the hands of ISPs, employers or tech platforms may create market-based incentives for security. For example, some users may avoid wanting to have their content monitored and are willing to assume the security risk that implies. Others may willingly subject content to inspection to minimize the security risk. Yet these market-based incentives may be thwarted by market concentration (particularly in the case of ISPs and tech platforms) and the quasi-public nature of ISPs in some countries.
- In general, organizations in a position to monitor traffic must balance the risk of overly intrusive monitoring that violates privacy with the potentially heightened security that could be guaranteed.

Two important countervailing technological trends impact the extent to which different actors are able to monitor internet traffic and enforce security protocols:

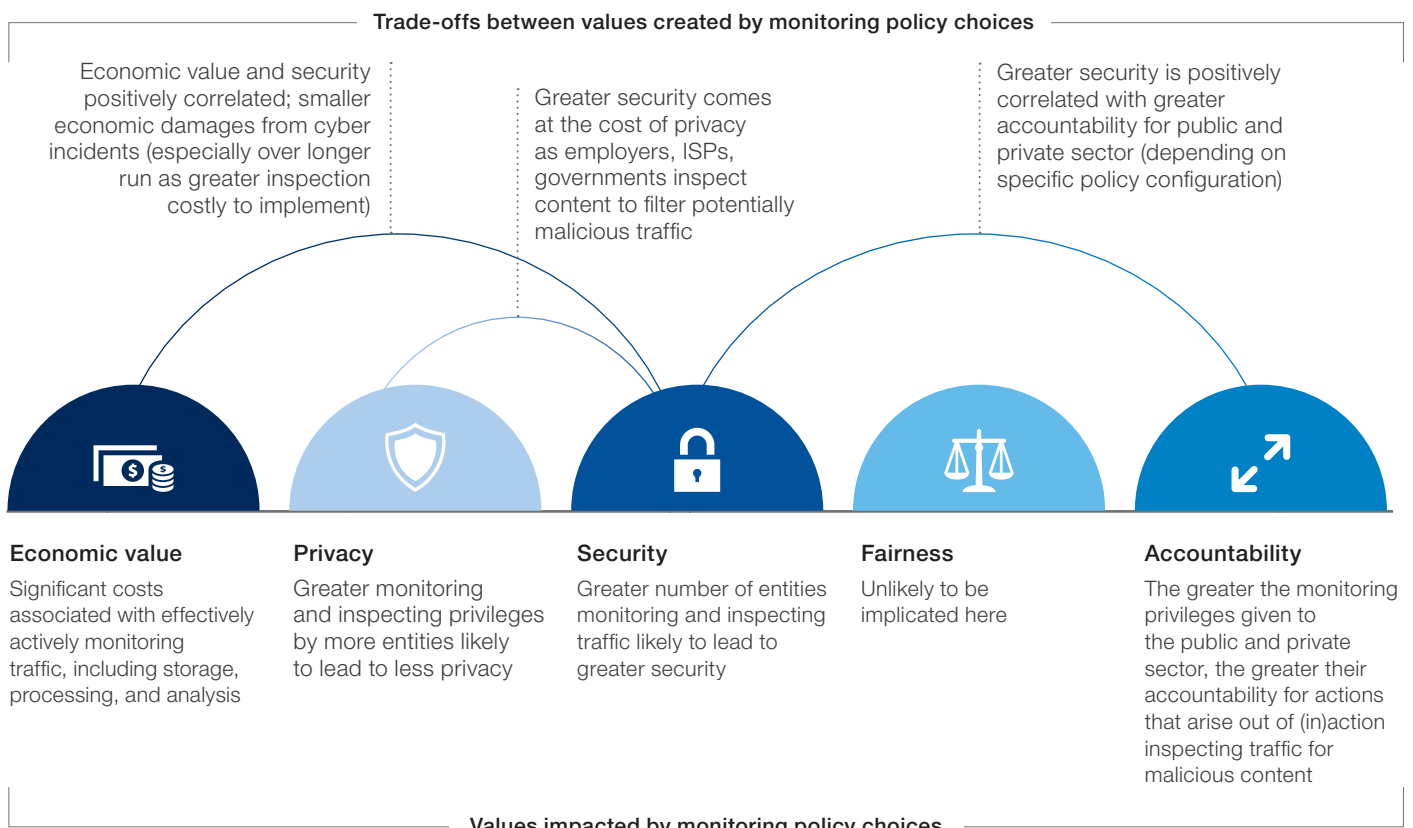
- Increasing amounts of data flows are being encrypted by default, thereby stymieing the ability of government, ISPs, tech platforms and, to a lesser extent, employers from observing and filtering content, depending on product and context (e.g. data “in flow” vs data “at rest”) even if it were a priori permissible.

Policy model: Monitoring



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by monitoring policy





4.6 Monitoring

- However, in some contexts the increasing adoption and proliferation of technology allow greater inference of content based solely on metadata. One example: if that metadata reveals the video compression protocol, the size of the packet transmitted and the address accessed, so-called deep packet inspection statistical techniques adopted by some ISPs could reveal that the precise size of the encrypted data has the digital “fingerprint” of accessing a particular form of objectionable content provided by a known website.³⁵

Connecting policy to values

The value trade-offs implicated by enabling greater monitoring — whether by employers, government, ISPs or tech platforms — are to a large extent the same value trade-offs associated with weak encryption policy, whereby law enforcement has a mechanism to contravene encryption. However, one key difference, at least to date, has been that monitoring privileges are typically the province of the private sector. Consequently, whereas “backdoors” are likely to deteriorate trust and create security issues whose combined effect might reduce commerce, monitoring privileges by private-sector intermediaries have not yet deteriorated trust in ICT:

- Greater monitoring privileges may improve security, provided those privileges are not compromised by bad actors who misuse the information gained. Monitoring content will allow different ecosystem intermediaries to filter and inspect traffic for malicious content.
- Greater security will result in economic benefits (assuming trust is not compromised). Given that fewer individuals and organizations will fall victim to cyberattacks, increased monitoring should reduce the costs associated with cyberincidents.
- Monitoring policy impacts the accountability of the public and private sectors. Given the ability to inspect content, both the public and private sectors will be enabled to use technology (of which deep-packet inspection is just one example) that will allow them greater capabilities to provide security for end users. Consequently, greater monitoring privileges will be associated with greater accountability.
- Privacy is also impacted by choices in monitoring policy. Greater monitoring, particularly of content, will reduce privacy.

4.7 Assigning national information security roles

Definitions

Robustness — being capable of performing without failure under a wide range of conditions. Cybersecurity measures that promote robustness can be classified into several main efforts, including: 1) organizational processes; 2) technical steps, such as network segmentation, user privilege policy, access control, data encryption and authentication mechanisms; and 3) procedures focused on the human factor, such as training

Resilience — the capability to detect threats, prevent their infiltration or at least confine their expansion, manage their effects and deny their recurrence; the notion of adaptability is at the core of resilience, as is being able to continue ordinary operations

Defence — the capacity to disrupt cyberattacks by focusing on the human factor behind them through national operational defence capabilities³⁶

Policy model

The question of establishing national cybergovernance is fundamental to ensuring security. What are the roles, responsibilities and capabilities that should be expected of the public and private sectors? Leaving aside issues of liability and insurance, it is important to establish clear roles and responsibilities for security within a country. When security is vaguely defined as “everyone’s problem”, or perhaps more dismissively “someone else’s problem”, in practice it is no one’s problem.

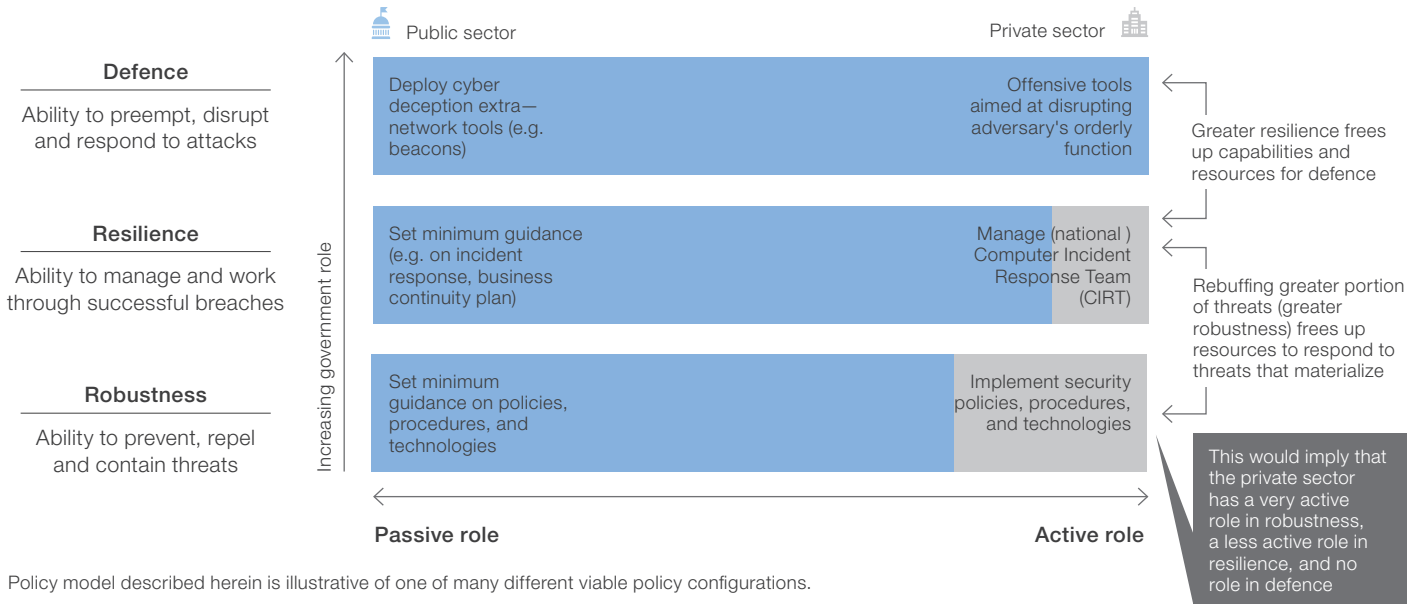
In contrast to policy in other spheres, organizations (rather than individuals) are the unit of analysis and action for this sphere. Security governance should be framed in collective terms (similar to threat intelligence) to create collective immunity. Furthermore, it should be framed in collaborative terms to leverage the private sector’s decentralized contextual knowledge and the public sector’s broad view of the threat landscape and considerable resources.

A three-layer additive framework can be used to help align roles and responsibilities with distinct security capabilities around software and the assets that software impacts, for which a national government should assign responsibilities. One preliminary determinant of how policy should be implemented in any given national context is ownership and responsibility of data. The World Economic Forum explored how to assign responsibility around personal data in an earlier publication.³⁷ The three layers of national cybergovernance are:

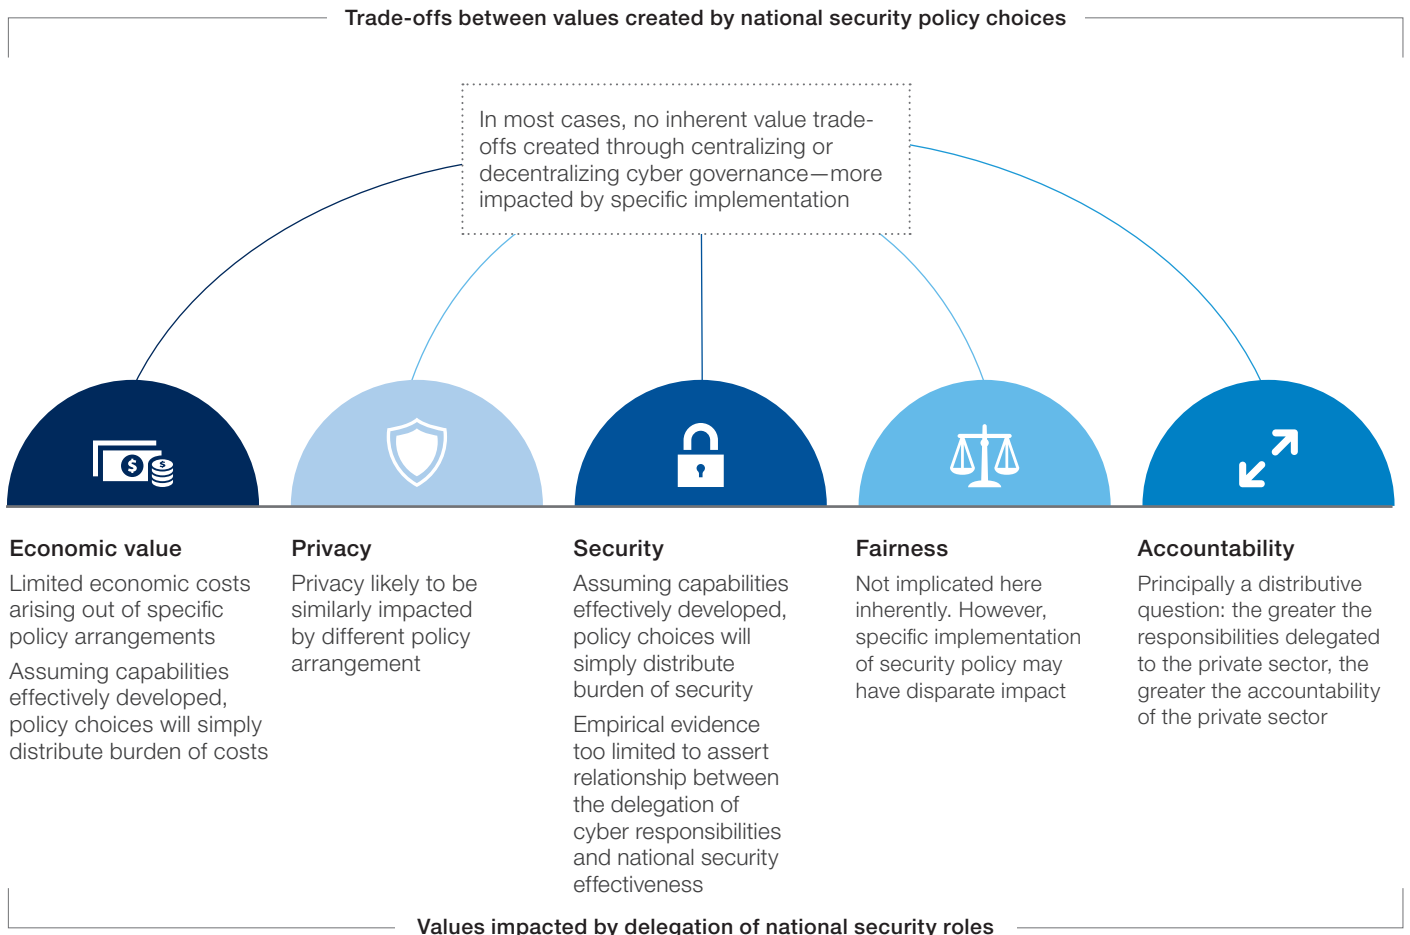
1. Robustness, which can be understood as the ability to prevent, repel and contain threats. In practice, this would consist of organizational and technological measures to prevent cyberincursions. For many countries, the first layer of defence is a responsibility delegated principally to the private sector, with the public sector providing guidance on standards and minimum policies, procedures and technologies.
2. Resilience, which is the ability to function during and after successful cyberincursions. One of many capabilities that helps promote resilience is a Computer Security Incident Response Team (CSIRT). Governments typically have taken a greater role in providing resilience, although large private organizations are also able to field capabilities that promote resilience. The upfront costs associated with resilience capabilities are high while the return for most private-sector organizations is episodic. Furthermore, such costs are highly scalable; while it is (usually) not feasible for any single organization to develop emergency response capabilities equivalent to a CSIRT, when those capabilities exist, an individual organization’s ability to respond to a breach is heightened by potential recourse to a CSIRT. Another important resilience capability is also business continuity planning, ensuring that organizations are prepared to manage through incidents.
3. Defence, the ability to pre-empt, disrupt and respond to cyberattacks. In contrast to other governance capabilities, which are fundamentally introspective, defence is focused on the originating source of cyberattacks. Again, defence is a role more naturally suited for government, given the exercise of sovereign responsibilities, laws and regulations related to intentionally doing harm to another individual or entity, and the economic profile of developing defence capabilities. Developing and maintaining defence capabilities is resource intensive, while the benefits are diffuse and over longer periods of time (e.g. deterrence). To be sure, there is considerable debate about the extent to which the private sector should be allowed to act to defend itself (which is addressed separately under point 4.12 “Active Defence”).
4. Each capability strengthens the others. Greater robustness means that governments will be required to deploy resilience less frequently. And greater resilience implies greater capacity can be dedicated towards defence. Alternatively, greater robustness and resilience capabilities might necessitate less of an investment in defence.

4.7 Assigning national information security roles

Policy model: assigning national information security roles



Key values trade-offs created by assignment of national security roles



Significant trade-offs are associated with delegating these capabilities to the public sector vs the private sector. The principal trade-off is that associated with governance centralization:

- The more robustness is a decentralized responsibility of the private sector, the greater the risk that friction in cross-organization intelligence sharing impedes security. To take a simplified example of disseminating a known suspicious URL, in a more centrally managed (through the public sector) robustness capability, it would be easier to ensure that less traffic is directed towards that URL. Alternatively, a more centralized security posture trades the ease of management for agility. Every organization faces unique threat vectors that may be overlooked if the public sector takes a more active role.
- A similar risk is associated with resilience capabilities. As resilience becomes centralized, the contextual knowledge of an organization’s specific network topography is less accessible for the public sector, which means that efforts to respond to cyberemergencies may be slower. Alternatively, as resilience becomes the responsibility of the private sector, inordinate (duplicative) costs may be borne by organizations (unless those capabilities are outsourced).
- The risks associated with defence capabilities are somewhat analogous to the trade-offs involved in allowing the private sector to take a leading role in attribution (as identifying adversaries is core to defence capabilities). The more defence is delegated to the private sector, the greater the risk of potentially significant collateral consequences (e.g. the private-sector hack back of an alleged public-sector adversary).

One general trade-off that should be considered across all the capabilities is the extent to which guidance, both from the government to organizations and from organizations to their constituents, is fully specified. The greater a given control is specified (e.g. an organization must have an antivirus solution installed on all endpoints), the greater the risk that organizations “solve for” security through compliance rather than the regulatory objective (e.g. secure endpoints). However, from the perspective of managing the regulatory apparatus, determining and verifying compliance from the perspective of government is simpler and less

expensive than understanding if regulatory objectives are being achieved. While this is generally true of many regulatory efforts, in the context of cybersecurity, the increasing heterogeneity of specifications made by well-intentioned actors has created a proliferation of requirements with varying effectiveness.

Case study: Waking Shark, United Kingdom

Since 2011, UK financial authorities have conducted cyber stress tests, Waking Shark I and II, on the financial sector to prepare and assess its readiness to increasingly severe cyberattacks. In contrast to typical “red team” exercises, in which external teams probe and hack a given company, “Waking Shark” is an industry-wide government-facilitated exercise to test financial infrastructure generally. A number of valuable lessons can be learned from the British experience:

- There are no substitutes for experience, but cyberexercises are one of the best ways to test an organization’s capacity for robustness and resilience and to challenge leaders to deal with the question of defence.
- Industry-wide exercises facilitated by the public sector are a unique form of collaboration. Conducting such a test would be difficult in a purely private context — organizations would be loath to participate and risk reputational damage with their counterparties, competitors and customers.
- The exercises have focused on the health of market infrastructure rather than on any individual participants. Not only does this focus allay the concerns of individual participants, it reveals a nuanced understanding of the connected nature of cyber-risk. Market infrastructure can only be as resilient as the weakest contributing link.

4.7 Assigning national information security roles

Case study: Cyber Star, Singapore

In 2017, the Cyber Security Agency of Singapore (CSA) conducted an exercise covering all 11 designated “critical information infrastructure” sectors in Singapore, in a whole-of-government effort to test Singapore’s cyberincident management and emergency response plans. The exercise comprised a series of complex scenario-planning sessions, workshops and table-top discussions, covering different types of cyberattacks targeting essential services, including web defacement, widespread data exfiltration malware infections, ransomware hits, distributed denial of services attacks and cyberphysical attacks. Participants also developed and tested their incident management and remediation plans in response to these simulated attacks.

To understand why such an expansive exercise is useful, it may be helpful to use the example of the financial sector, typically deemed critical in most countries. The ability of the financial sector (or any sector) to withstand a cyberattack is deeply premised — sometimes unquestioningly — on the availability of adjacent sectors. How many banks can withstand a cyberattack when their ISP is besieged? To what extent can market infrastructure absorb the deterioration of power generation capabilities through a cyberattack on industrial control systems? If transportation infrastructure is constrained (e.g. metropolitan transit no longer operates), who will staff the security operation centres of financial institutions? The scenarios and potential linkages are numerous and without some level of planned and exercised coordination, it is difficult to imagine how resilience capabilities will be maintained during a cyberincident.

Connecting policy to values:

The inherent value trade-offs created by choices in cybergovernance defining how information security roles and responsibilities will be delegated are highly context dependent. However, a few efforts across economic value, privacy, security and accountability reveal general themes:

- Some capabilities have the profile of a pure public good (in the classic economics sense): their consumption is non-rivalrous and non-excludable. Deterrence arising as a consequence of defence capabilities would be one such example. As such, the economic value of the public sector providing this capability is likely to be greater. Other capabilities have a more mixed profile.
- A given polity’s understanding of privacy is highly dependent on the context and may vary depending on how responsibilities are delegated. In some contexts, a more active role for the public sector in providing robustness capabilities may be perceived as a diminution of privacy, whereas in others the ability to access sensitive information would raise symmetric privacy concerns regardless of whether it is led by the private or public sector.
- Security can be achieved through a variety of assignments of roles and responsibilities — little empirical evidence suggests that a more or less centralized role for the public sector necessarily results in greater security. However, it is worth noting that the degree of centralization of cybergovernance differentially impacts the security risks that are mitigated. A greater degree of centralization, in which the public sector has a more active role across robustness, resilience and defence, is likely to be more effective at addressing coordinated and broad threats. Contrariwise, a greater degree of decentralization is likely to be more effective at addressing diffuse and heterogeneous threats.
- Governance decisions will necessarily impact the accountability of the public and private sectors; governance choices will principally distribute that accountability. However, given that private-sector organizations are typically the providers of ICT, for some capabilities (e.g. resilience) the extent to which private-sector accountability diminishes if the public sector takes a more active role is limited.

4.8 Encryption

Definitions

Encryption — the cryptographic transformation of data (called "plaintext") into a form (called "cipher text") that conceals the data's original meaning to prevent it from being known or used

Strong encryption — encryption that cannot be decrypted through reasonably accessible computational methods or algorithmic flaws

Weak encryption — encryption that can be decrypted through reasonably accessible computational methods or algorithmic flaws; additionally, also considered weak is strong encryption that has a built-in bypass capability (commonly referred to as a "backdoor")

End-to-end encryption — a system of communication in which the only people who can read the messages are those who are communicating; no eavesdropper can access the cryptographic keys needed to decrypt the conversation — not even a company that runs the messaging service³⁸

Policy model

Encryption is a fundamental technology for security. The key policy question for understanding how to treat encryption is: who should be able to access sensitive data and communications? Encryption is necessary to ensure that sensitive data and communication are not accessed by bad actors. However, encryption can also be used by bad actors to shield communications from law enforcement. In the last few years, encryption has become increasingly salient as the private sector has invested in differentiation on the basis of user-friendly encryption to secure increasing amounts of personal (sensitive) data. On the other hand, some policy-makers increasingly insist on weaker encryption. On this particular policy topic, minimal opportunity for a middle ground exists. An encryption algorithm either obfuscates data or it does not. And algorithms cannot divine the intentions of those seeking to circumvent them.

To help frame the policy discussion for encryption, it is helpful to think about policy on two (related) axes: who has access to encryption (e.g. the public or private sector) and what type of encryption do they have access to?


- Two analytically helpful (though not necessarily technological) types of encryption exist: weak and strong. In some circumstances, one entity's strong encryption might be considered weak by another. For example, encryption that may be used by consumers may not be sufficiently strong for defence ministries. However, the policy-relevant deliberation is not about technological standards but about the choice of whether to allow access to data by an entity other than the user regardless of the underlying technology.
- For the purposes of encryption, two entities are fundamental: the public sector and the private sector. Again, within the public and private sectors, different subgroups might utilize unique technologies around encryption but the question is about access and its purpose.

Requiring differential encryption (particularly for the private sector) has significant risks and benefits:

- If the private sector is required to use weak encryption, then bad actors may potentially obtain access to customer data. The greater the amount of customer data that companies are allowed to capture, the greater the potential damage associated with bad actors accessing this data.
- Alternatively, if the private sector is allowed to use strong encryption, law enforcement may be hampered in its efforts to access data relevant to preventing crime or investigating its aftermath. Strong encryption also stymies intelligence agencies in the collection of data.
- A policy of weak encryption for the private sector may be unstable in the long run when coupled with allowing private-sector access to greater amounts of personalized data as the cost of bad actors defeating encryption may become greater than the value of thwarted criminal activity. Requiring private-sector vendors to develop encryption workarounds may also impose non-trivial costs, not only in terms of customer risk but in terms of software development and engineering costs.

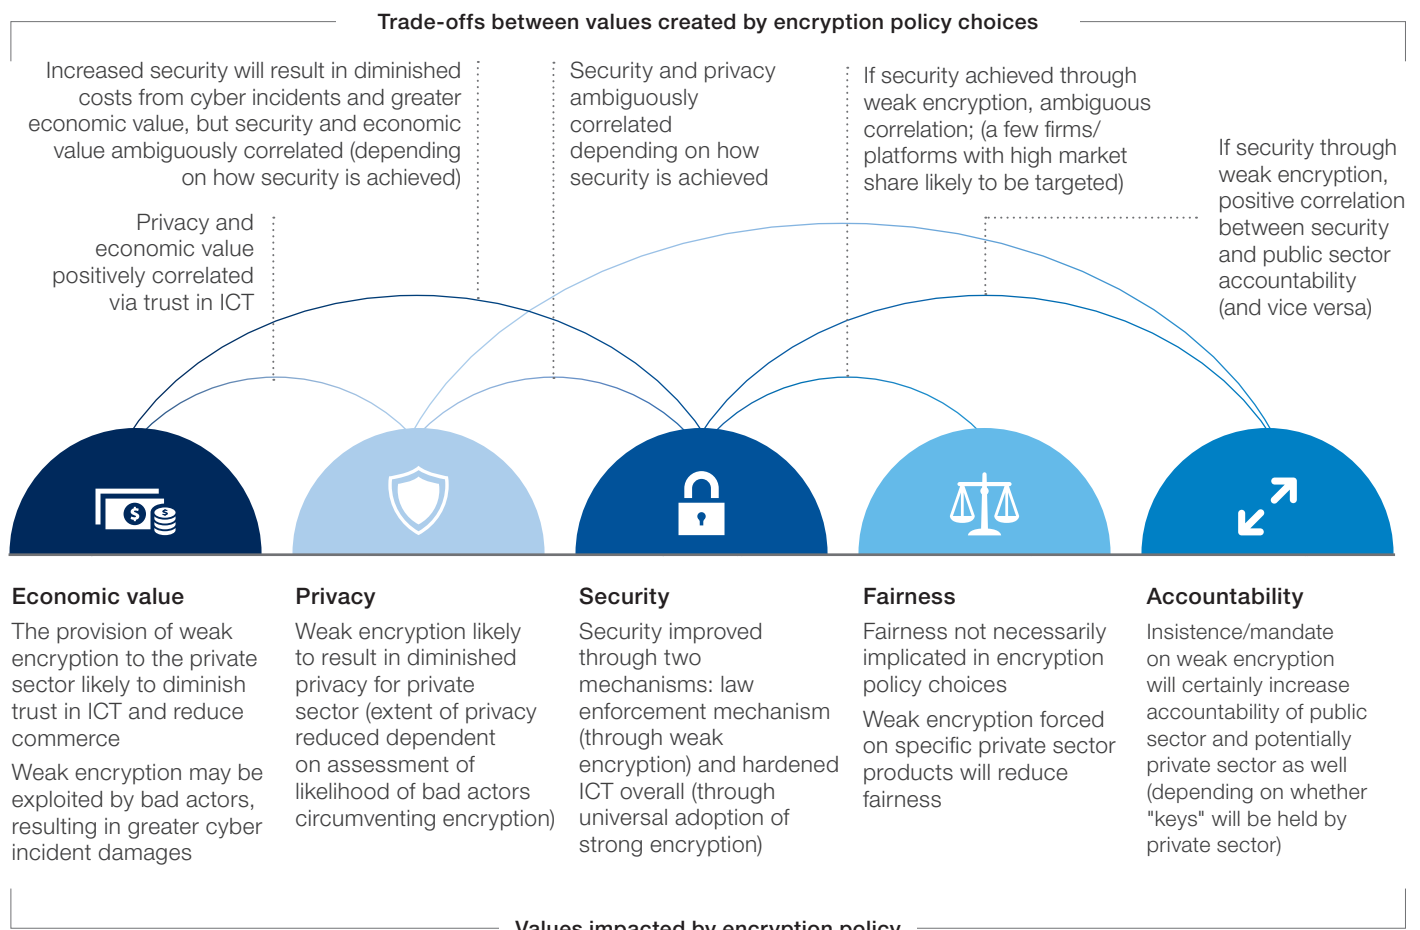
4.8 Encryption

Policy model: Encryption



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by encryption policy choices

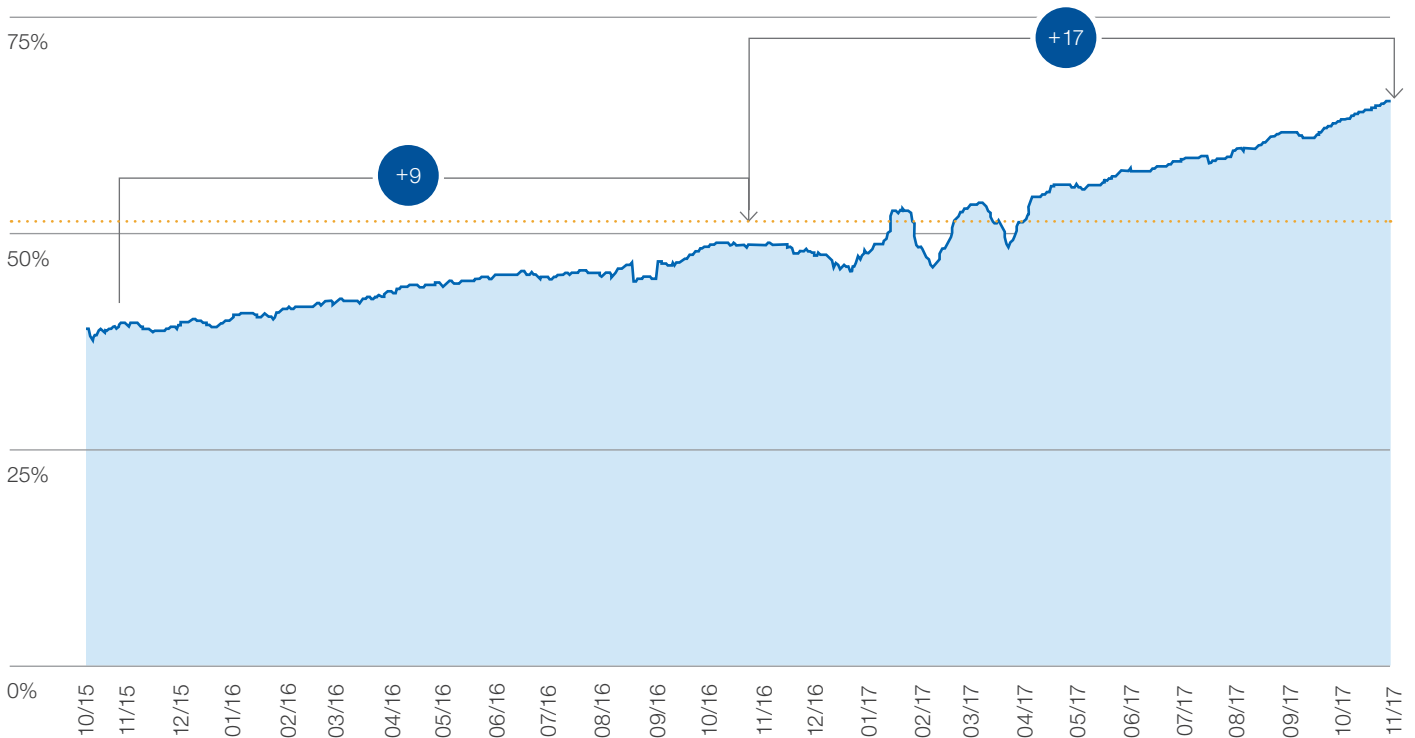


To understand the concerns and risks around weak encryption, it may be helpful to reason through a commonly debated example: if policy-makers insist on a way to bypass a popular messaging application — regardless of whether that application can currently support such a bypass — bad actors will try to move to other applications to mask communications. If policy-makers then insist on achieving access to communication at a more fundamental technical level (e.g. at the level of an operating system),

so that bad actors have no choice but to (at least in the short run) risk exposing their communications, other bad actors in cyberspace will have even greater incentives to pierce that encryption. After all, operating-system-level privileges are valuable to law enforcement/intelligence for the same reason as they are to adversaries. This logic can continue down the technical stack but, in general, the more inescapable the bypass capability sought, the more attractive that bypass becomes to adversaries.

Increasing adoption of encryption for internet traffic

Percentage of page loads over HTTPS



Source: World Economic Forum; BCG analysis.

14-day moving average; <https://letsencrypt.org/stats/#percent-pageloads>

HTTPS is commonly used encryption protocol for securing web traffic. By contrast, HTTP is not encrypted.

Case study: Encryption and business model disruption

Some commentators have encouraged the adoption of end-to-end encryption to ensure that data-at-rest and data-in-transit remain secure from unauthorized access and disruption.³⁹ Responding to some policy impetus (e.g. General Data Protection Regulation in Europe), companies are increasingly implementing end-to-end encryption. However, it is important to acknowledge that end-to-end encryption threatens business models premised on monetizing individual-specific attributes or using individual-specific data for advanced analytics (including personalized AI and machine learning). In adopting end-to-end encryption, companies limit the ability to inspect communications. In so doing, companies limit the inferences they are capable of making regarding an individual — whether that individual is likely young or old, male or female, etc.

As a consequence, the ability to then sell an adjacent service (e.g. advertising or an AI-based service) targeting individuals based on revealed attributes is greatly diminished. In the case of advertising, the explanation is reasonably straightforward: the ability to target and measure the impact of an ad is paramount for marketing teams to articulate a value proposition to negotiate for budgetary authority. In the case of advanced analytics, the impact of encryption of those business models is a bit more subtle. In general, advanced analytics require the aggregation of both data and computing typically limited to accessing a cloud resource. However, some companies have experimented with using mathematical models capable of inference that never leaves an endpoint or is obfuscated when in transit to cloud resources, such that data remains anonymized and encrypted. Nonetheless, the trade-off is clear: encryption obfuscates data that could otherwise form the basis of the richer inference underlying personalized AI/machine learning-based services.

Case study: Quantum computing and encryption

Recently, companies have begun to commercialize access to quantum computing. In light of this access, some commentators have raised concerns about the ability to encrypt data in light of these new computational techniques. These concerns are somewhat alarmist; while it is true that the current mathematical algorithms underlying much of the encryption used by the public and private sectors would be vulnerable to these new computing techniques, already efforts are in place to develop new algorithms to thwart quantum computing. The U.S. National Institute of Standards and Technology (NIST) has begun developing so-called “post-quantum” cryptographic techniques.⁴⁰

Connecting policy to values

Encryption policy choices sharply implicate a number of values and, depending on policy choice, create key trade-offs between these values. The value trade-offs surfaced by encryption policy are very similar to those raised by zero-day policy. It would be inconsistent, for example, to argue for pervasive stockpiling of zero-days while insisting on strong encryption for the private sector — vulnerabilities in encryption make it weaker:

- Security may theoretically be improved in two scenarios: one in which the private sector has weak encryption or strong encryption. One line of thinking, more associated with law enforcement, is that governments can provide greater security for citizens and firms by accessing communications that may be used by criminal elements in a weak encryption policy.
- Another line of thinking suggests that strong encryption is likely to promote greater security, such that bad actors do not discover and exploit backdoors to encryption against a country’s citizens and firms. A weak encryption policy is more likely to mitigate the risk of coordinated and broad threats, as law enforcement access will presumably deter would-be conspirators and facilitate rapid criminal response. A strong encryption policy is more likely to mitigate the risk of bad actors exploiting sensitive information.
- The economic value associated with different encryption policy scenarios is a function of a few effects. For instance, greater security achieved through weak encryption and strong encryption is associated with fewer damages arising from cyberincidents. However, in the case of weak encryption, this must be weighed against the costs of these same backdoors being used against a given country’s citizens and firms, as well. Additionally, some observers have noted that actions deteriorating trust in ICT create substantial intangible costs in terms of diminished ICT adoption.
- Privacy is also impacted by choices in encryption policy. In a weak encryption policy, the improvement in security is premised on decreased privacy. While, in most cases, presumably decreased privacy will be limited to suspected criminals, the risk is that the confidentiality of non-adversaries will also be compromised.
- Encryption policy impacts the accountability of both the public and private sectors. It is incumbent on the private sector to adopt sufficiently strong encryption to thwart adversaries. However, if the private sector is mandated to use weak encryption, the public sector has greater accountability to ensure that backdoors remain undiscovered and that increased access to communications is closely monitored and also productively used by law enforcement.

4.9 Cross-border data flows

Definitions

Data sovereignty — the concept that information is subject to the laws of the country in which it is located; many of the current concerns that surround data sovereignty relate to enforcing privacy regulations and preventing data that is stored in a foreign country from being subpoenaed by the host country's government

Data localization — barriers to cross-border data flows, such as data-residency requirements that confine data within a country's borders

Policy model

Despite the best efforts of early internet pioneers, cyberspace has become a domain subject to nation-state sovereignty. The increasing imposition of sovereignty on cyberspace is a natural corollary to the internet's growing implications for a nation's well-being and security.

However, discussion of the relationship between national sovereignty and cyberspace has often been incomplete, focusing only on the traditional question of data sovereignty: limiting the transit of personal data across national borders.

To appropriately account for the costs and benefits of sovereignty in cyberspace, it is necessary to take a wider view that evaluates both data egress (e.g. data sovereignty, export controls on cryptographic protocols) and data ingress (e.g. content restrictions):

- **Data egress** — Governments have undertaken measures to limit the data that leaves national borders. One salient example of these measures are growing data localization efforts. In the wake of concerns that citizens' personal information may be surveilled or monetized by corporations, governments have begun to limit the extent to which personal data can cross national borders. As a consequence, in a "cloud-first" world where companies are increasingly creating software offerings premised on access to a centralized pool of resources and applications, data egress limitations are resulting in companies investing in more localized data centres. However, it is unclear whether data localization efforts effectively constrain a foreign government's access to data on a given host country's citizens.⁴¹

- **Data ingress** — A more recent trend is policy-making around limiting the data that enters a given country. Establishing national barriers to the spread of information is not a new phenomenon but the application of these barriers to cyberspace is novel. While a number of salient examples exist, the most visible expressions of limitations on data ingress are content-restrictions that various national governments have established to limit citizens' access to certain websites or content.

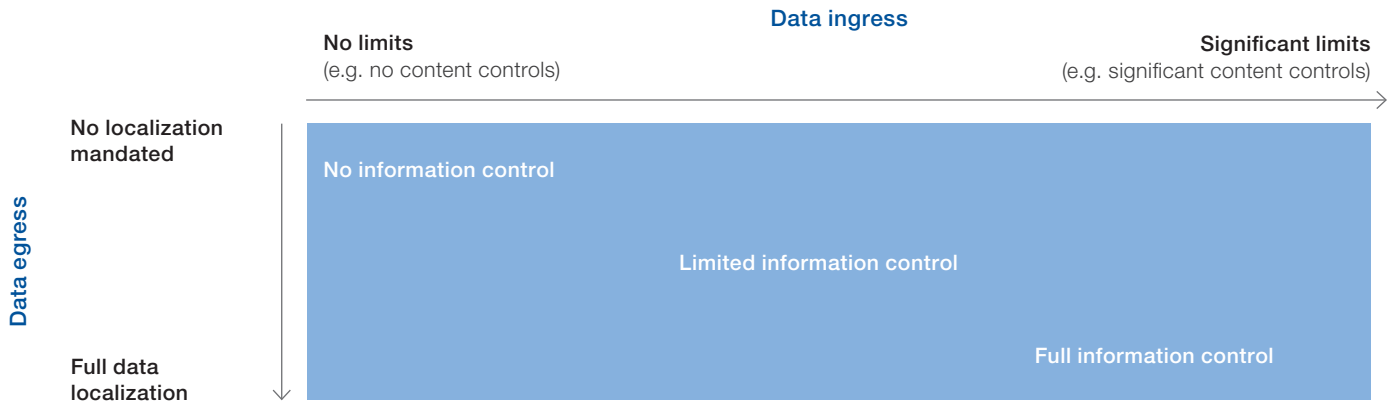
Increasing efforts to exercise sovereignty in cyberspace have significant risks and benefits:

- Whether it is expressed in terms of limitations on data egress or data ingress, in general, efforts to increase information control impose additional costs on the users of internet services. For example, if a given country undertakes data localization efforts, cloud-based services will be more expensive to deliver and, in some cases, offered with delays. After all, service providers must amortize the additional cost associated with building or accessing additional relatively expensive data centres.
- Additionally, data localization efforts have a mixed impact on security. While localization legislation may be embraced as an opportunity to set clear policy on security generally, the proliferation of physical data centres (beyond those needed for redundancy) is a security risk because there are more physical targets.⁴²

Policy-makers may consider borrowing from the relatively well-developed intellectual framework of trade economics in considering questions of cross-border data flows. For example, one counter-intuitive finding of trade theory is that an import tax may be effectively borne by exporters. Similarly, efforts to limit "importing" data (e.g. content restrictions) place a heavy burden on data "exporters". For instance, engineers unable to freely query and refer to global experience in software development are likely to face greater difficulties developing software.

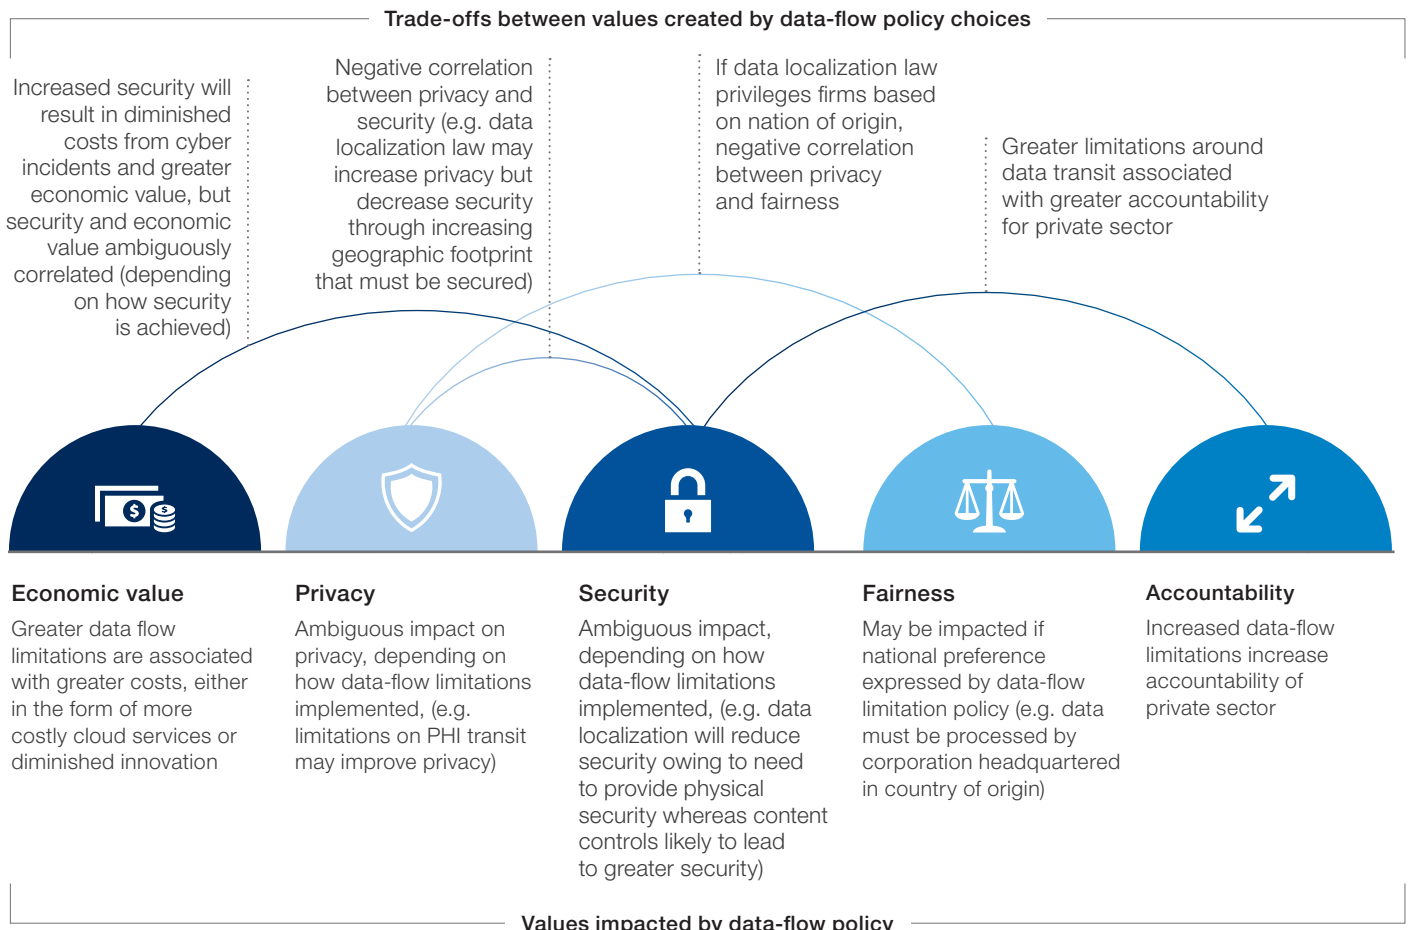
4.9 Cross-border data flows

Policy model: Cross-border data flows



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by data flow policy choices



Case study: The economics of data centers

Within a given country and also in an international context, the economics of data centres — the physical linchpins for cloud resources — are commonly misunderstood. In general, data localization (and the subsequent reshuffling of data centres) imposes much greater costs than benefits for any subnational locale or country:

- Within a given country, there is often intense competition for the promise of enormous investment by companies building data centers, typically through tax incentives. But the capital expenditures associated with a datacenter result in little long-term employment. Indeed, that is in some sense the motivating principle of a datacenter—how to build cloud resources with the lowest recurring operational costs whether it is electricity or people-related costs. The canonical example illustrating these dynamics is a \$1B data center built by Apple in North Carolina that created “only” 50 jobs.⁴³
- Some policy-makers argue in favour of data localization efforts on the basis of the economic benefit of bringing data centre construction to a given locale. However, the economics of an incremental data centre in a new locale are similarly self-defeating. While there are limited concentrated benefits associated with the construction of a data centre (per the Apple example), the costs associated with the localized provision of cloud services are diffuse and non-trivial. A recent Information Technology & Innovation Foundation report benchmarked these effects using memory allocated for storage and found that data localization greatly increased costs for local companies: between 10.5% and 62.5% more for some cloud-computing services.⁴⁴ These increased costs are bounded by the availability of alternatives (e.g. a company builds their own private data centre instead of relying on cloud resources).
- Data localization costs are not only imposed on users within a given country but also internationally. The providers of cloud resources, despite the increased costs noted above, also amortize some of the costs imposed by a given locale across the entirety of the customer base.

Connecting policy to values

There are few inherent value trade-offs associated with data flow policy choices in the abstract — a number of different polities have implemented and administered data flow limitations with differing effectiveness and impact depending on the national context.

- Increased cross-border data flow limitations may improve security insofar as they codify and organize national policy on personal data. In other words, the limitation itself is unlikely to provide security (given the exacting security controls multinational cloud service providers already adopt) outside of policy clarification. That said, data flow limitations, which amount to a mandate to build physical data centres in a given locale, may reduce security depending on the physical security of those data centres and the trustworthiness of ancillary network infrastructure.
- Increased data flow limitations will almost certainly increase costs to a greater extent than the security incident damages averted owing to greater security. Data flow limitations have significant direct costs (e.g. more expensive cloud resources) and indirect costs (e.g. decreased cloud adoption and slower innovation).
- The impact of data flow limitations on privacy and fairness is ambiguous. For example, increased limitations on the handling and processing of personal health information (PHI) may improve privacy. On the other hand, limitations on the content an individual can access may intrude on an individual’s privacy. Data flow limitations may be entirely fair and neutral (e.g. all cloud providers must adopt certain controls for the transit of personal financial data). Alternatively, data flow limitations may unfairly privilege companies based on national origin (e.g. data in a given locale must be processed by a corporation headquartered in the country of national origin).
- Increased limits on cross-border data flows will almost always increase the accountability of the private sector. Administering data flow limitations will be a private-sector-led effort in most contexts, and as such it will be the responsibility of that sector to ensure that specific limitations are affected.

4.10 Notification requirements

Definitions

Personally identifiable information (PII) — any data that could potentially identify a specific individual; any information that can be used to distinguish one person from another and can be used to de-anonymize anonymous data can be considered PII. Breach notification laws typically focus on notifying the public when PII might have been exposed to unauthorized individuals, particularly in the context of financial or medical information

Breach — an incident in which sensitive, protected or confidential data has potentially been viewed, stolen or used by an individual unauthorized to do so; data breaches may involve PHI, PII, trade secrets or intellectual property⁴⁵

Policy model

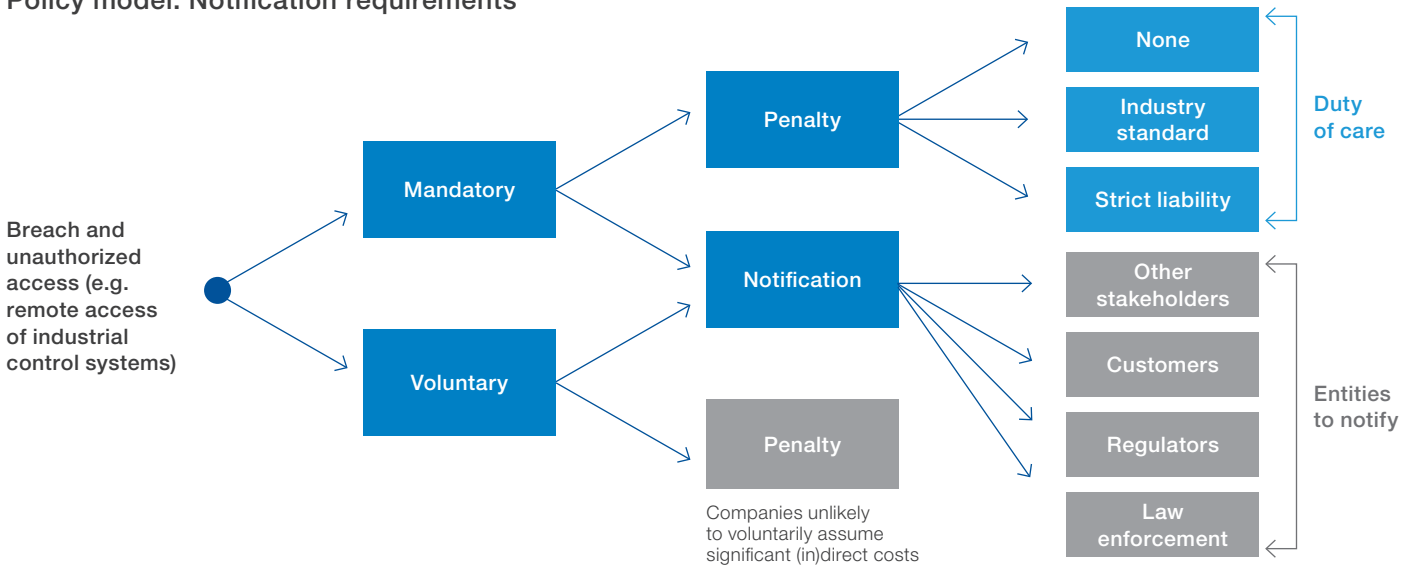
As more companies are successfully attacked in cyberspace, policy-makers are trying to develop procedures around informing customers, regulators, citizens, investors and other affected stakeholders when sensitive data is potentially compromised.

Two major axes define the contours of breach notification policy:

1. First, when are companies mandated to report a cyberincident? Or, is notification a voluntary disclosure? Which stakeholders should be notified? For example, policy-makers could fashion a hierarchy of notification whereby it would be mandatory to notify law enforcement but voluntary to notify other stakeholders.

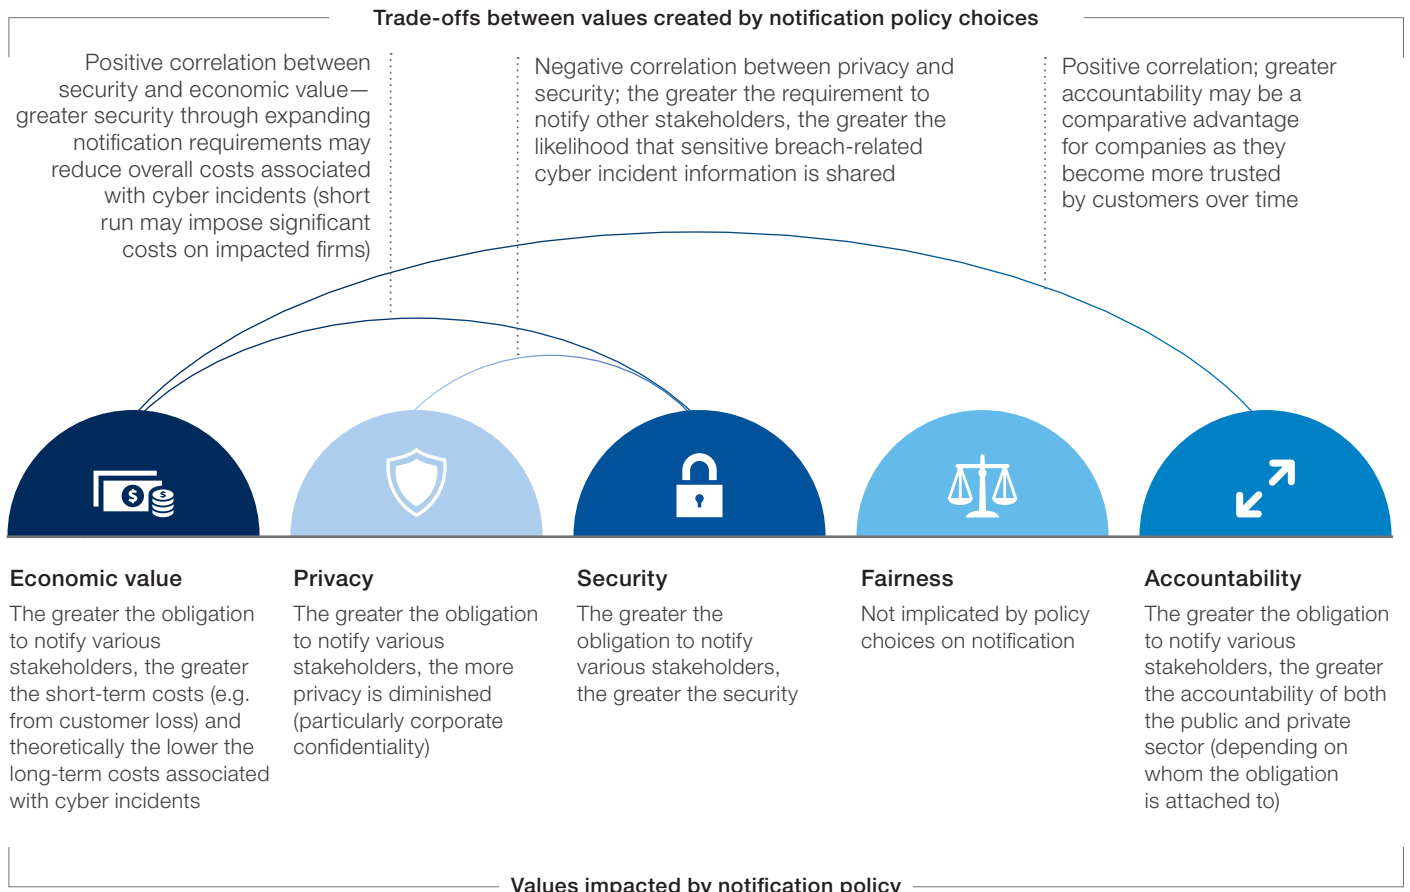
2. Second, what form of sanction is attached to the breach, itself?
- Should companies pay penalties? On what basis should those penalties be levied? Policy-makers may elect between three broad levels of care that might trigger penalties. The first level, and least stringent, would be to attach “no penalty”, so as to avoid punishing companies victim to an attack. The second would be to penalize companies if they did not maintain a level of care (e.g. consistent with industry standards). The third, and most demanding, would be a policy of strict liability. Companies would be penalized regardless of the duty of care they exercised.
3. Additionally, a few important additional questions must be asked in crafting policy:
- How long should companies have before they disclose a breach? To whom? For policy-makers, the trade-off they should be assessing to determine an appropriate amount of time before a company must report a breach is the following: would cyberincident damages be reduced to a greater extent by allowing a company time to manage an organized response or by allowing affected individuals to act earlier in a decentralized fashion? One additional consideration is that attempts to set national policy may be thwarted by international actions. To take a simplified example, a breach notification law with a 10-day window in one country will be effectively nullified by a breach notification law with a three-day window in another jurisdiction. Put differently, the lower common denominator will prevail. Additionally, within a country, specific enumerated time limits may create their own issues; cyberincidents differ and the extent to which a given stakeholder would benefit from knowledge of an incident by a given point in time will also materially vary.
 - How should companies notify relevant stakeholders, especially given the increasing frequency of breaches? These questions become especially salient when trying to combine notification with advisory measures for consumers on how to mitigate the damages caused by a breach.

Policy model: Notification requirements



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by notification policy choices



4.10 Notification requirements

For each of these policy choices, significant risks and benefits affect the incentives to invest in security and the resultant costs of cybercrime:

- If companies are required to report a breach, then investment in security will increase to avoid either embarrassing publicity or regulatory penalties.
- If companies are required to notify stakeholders and the public at large, they may also undertake additional investment in security to avoid the customer run-off away from an insecure business. One important consideration, however, is that customers and consumers, in particular, are becoming increasingly inured to breach notifications (otherwise known as “data breach fatigue”).
- If companies are required to pay penalties, particularly if these penalties are meaningfully additive to the expected outcomes associated with negative market sentiment, companies will invest still more in security. A policy regime that then attached strict liability to a breach would result in enormous increased investment in security.
- In all of these cases, increasing sanctions will create the classic trade-offs associated with security investment, including diminished opportunities to invest in other parts of a business, or a more general increase in operating costs that might be passed to users.

Connecting policy to values

Notification policy has important implications for a number of values principally animated by the extent to which such policy drives accountability:

- Increased notification requirements and breach-related penalties will increase accountability for organizations in the private sector. In the short term, increased notification requirements are likely to lead to greater costs for these organizations. Costs will increase as a consequence of regulatory penalties, consumer sentiment potentially shifting away from insecure companies, and the subsequent investment of those organizations in increased security controls. Over the longer run, the increased precautions taken by organizations should result in diminished costs associated with security incidents as their security improves. Furthermore, a competitive benefit may be realized by those organizations able to demonstrate more careful stewardship over sensitive data.
- The more notification policy becomes expansive and companies are required to report more incident-related data to various stakeholders, the more privacy is likely to be diminished (at least in the short term, until security improves such that user and corporate data is more likely to be safeguarded).

4.11 Duty of assistance

Definition

Critical infrastructure — systems and assets, physical or virtual, so vital that their incapacity or destruction would have a debilitating impact on national defence, economic security, public health and safety, or any combination of these matters

Policy model

One of the key questions confronting policy-makers is whether and how the government should draw upon public resources to assist an attacked private-sector organization. Clearly defining and circumscribing the public sector's duty to assist is an important and difficult policy topic. The public's resources to assist in a cyber emergency are finite and bounded. These valuable capabilities may exceed what is available in the private sector. As such, it is imperative to employ those resources judiciously and consistently. Within the context of the prior discussion on national information security roles, this policy model helps illustrate the key considerations to take into account when delegating responsibilities for resilience.

To help frame the policy discussion, it is helpful to think about the government's duty of assistance as contingent on two factors: the alleged identity of the adversary and the degree of risk. Additionally, it is worth noting that a duty can manifest in at least three behaviours by the government: no duty — in which the public sector is not obliged to offer assistance; an affirmative duty — in which the public sector is obliged to offer assistance without any obligation on the part of the organization to accept that assistance; and a mandate for an organization to accept public-sector assistance. To be sure, the provision of assistance may vary between national and subnational government based on national context, and may involve some form of public-private partnership:

- As the identity of the adversary triggering a duty of assistance may potentially impact a country's sovereign responsibilities, it is necessary for the government to be prepared to provide a more forceful response. For example, a government may choose to establish a legal duty for organizations to accept assistance, regardless of the potential damages observed, if they suspect that a nation-state actor was the originating source of the intrusion.

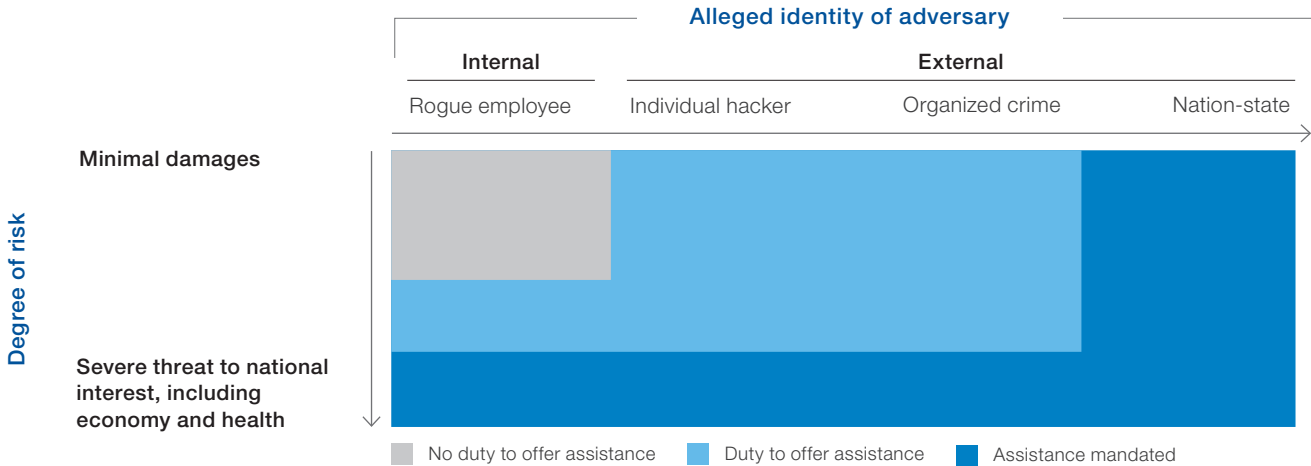
- The extensiveness of potential damages is a second key factor that should drive the forcefulness of the government's response. The consideration of risk (in the form of potential damages) here, as opposed to realized damages, reflects an important difference between the profile of cyberattacks versus other sorts of emergencies or disasters. Cyberdamages do not escalate linearly as a function of time — attacks moving at network speed may cause rapid stepwise increases in damages. One helpful example of the importance of thinking in terms of potential damage is the recent uptick in malware targeting critical infrastructure and, in particular, the electric grid.

Significant trade-offs are associated with assigning government a duty to assist “earlier” (in the case of smaller potential damages and less worrisome adversaries) or “later” (in the case of greater potential damages and more worrisome adversaries):

- The greater the scope of government duty, the greater the costs that must be borne to assist the private sector.
- One corollary to a more expansive government duty is that presumably adversaries will perceive a greater risk and thus be deterred.
- One additional consideration (discussed in greater depth in the context of liability thresholds) is that establishing a duty to assist may result in the private sector having diminished economic incentives to invest in its own emergency responsiveness.

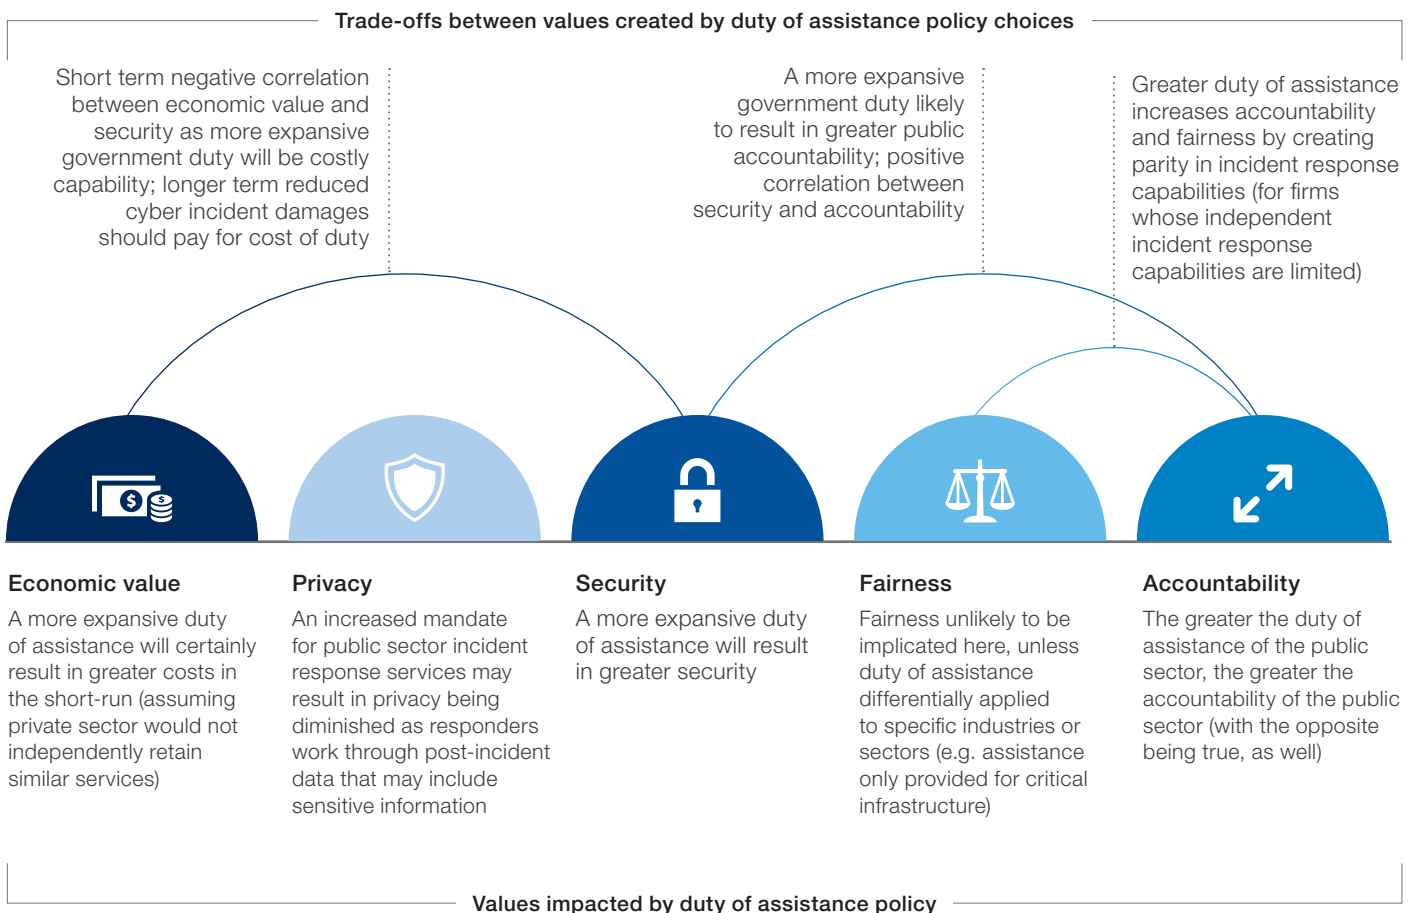
4.11 Duty of assistance

Policy model: Duty of assistance



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by duty of assistance policy choices



Case study: Defining critical infrastructure policy

Defining critical infrastructure is a key part of national policy that helps determine when a government's duty to assist ought to begin. The exercise of defining critical infrastructure is also important to help narrow the threat surface in scope and prioritize national assets for cybersecurity. However, defining critical infrastructure is fundamentally a context-specific exercise for any given country. In defining critical infrastructure, policy-makers should take into account a few important questions:

What attributes (and companies) qualify as critical infrastructure?

In the United States, for example, the Department of Homeland Security has outlined 16 sectors. Singapore has defined 11 critical information infrastructure sectors. Even within critical infrastructure, it may be valuable to prioritize sectors that may be more crucial to national security and economic well-being than others.

What elements of critical infrastructure policy are publicized?

Governments may choose to retain some level of ambiguity in disclosing how their duty of assistance is triggered. For example, most countries do not publicly disclose a list of companies that qualify as critical infrastructure. In so doing, governments are tacitly acknowledging that the value of "security by obscurity" is greater than the security potentially derived through deterring would-be adversaries.

How does government police the natural inclination to define critical infrastructure more broadly over time?

Governments may seek to define critical infrastructure more broadly to extend the umbrella of protection and to induce the private sector to upgrade its security. The private sector may also insist on inclusion as part of critical infrastructure. However, the extent to which government can serve its duty to assist is finite and mechanisms are essential to ensure that the duty is firmly circumscribed in scope and time.

One potentially useful policy analogy for critical infrastructure policy is financial regulation and crisis response. In the wake of the global financial crisis, countries went about balancing many of the aforementioned dynamics:

- Who should qualify for financial assistance? As part of the initial response to the financial crisis in the United States, the government mandated that financially "healthy" and "unhealthy" institutions accept capital injections to avoid the potential risk of "unhealthy" institutions being exposed to the financial strain of market participants' distrust. One analogue to critical infrastructure policy is not only offering assistance to attacked institutions but mandating it for a class of institutions to avoid scrutiny by adversaries. In the cyber context, assistance would have to be more tailored than a capital injection given the unique security requirements of each organization, but a similar policy intuition applies.
- Who should qualify for an outstanding government duty to assist? Again, in the United States, the government designated certain institutions to be "systemically important financial institutions". In so doing, the government attached certain demands to ensure that these institutions did not lose the economic incentive to mitigate their own risks (e.g. so-called "living wills" to ensure orderly bankruptcy). An analogue here could be to combine critical infrastructure status with exceedingly stringent security mandates (which already exist in some circumstances by virtue of certain critical infrastructure sectors generally being already heavily regulated).

4.11 Duty of assistance

Connecting policy to values

The extent to which the public sector's assistance is extended to the private sector raises sharp trade-offs between security, economic value, accountability and fairness:

- Provided the public sector has the effective capability, an increased duty of assistance to the private sector will likely result in greater security through a few key mechanisms. First, during a cyberincident, the public sector may provide effective incident response services. Immediately following a cyberincident, the public sector may provide resources and expertise allowing an organization to securely continue functioning. Finally, a greater duty of assistance may deter would-be adversaries.
- oWing to that greater security, the economic value of a greater duty of assistance would be positive in the long run. Many forms of cybercapabilities, particularly incident response, are exceedingly expensive to develop in terms of human capital. While this may result in some costs in the short run for a given country, the long-run benefit of effective incident response capabilities will outweigh the costs borne, at least in the present context where it is widely agreed that many organizations lack sufficient incident response capabilities.
- An increased duty of assistance for the public sector will also increase public sector accountability, but perhaps at the cost of the private sector's accountability.
- To the extent that the government extends its capabilities more broadly, such a policy would promote greater fairness — organizations that are less capable of responding to cyberincidents owing to context or resource constraints would be on a more level playing field with organizations that have developed or contracted for incident response capabilities.



4.12 Active defence

Definition

Active defence — a term that captures a spectrum of proactive cybersecurity measures that fall between traditional passive defence and offence (also sometimes colloquially known as hack back); active defence can fall under two general categories: first, technical interactions between a defender and an attacker, and second, operations that enable defenders to collect intelligence on threat actors and indicators on the internet, as well as other non-cyber policy tools (e.g. sanctions, indictments, trade remedies) that can modify the behaviour of malicious actors.⁴⁶

While some commentators have analytically differentiated hack back from active defence by the intention of the attacked organization (e.g. active defence refers to attempts to retrieve information whereas hack back refers to reciprocally inflicting damage on an alleged adversary), this report uses these terms interchangeably.⁴⁷

Policy model

An increasingly important question for lawmakers is: what limits should apply to active defence measures by a private organization? How should the government clearly circumscribe the technical measures a private-sector organization is empowered to use to respond to attacks? Such measures have created heated debate on both the technical and ethical fronts.


Active defence is a topic of controversy for practitioners, as certain common practices used to investigate and respond to potential intrusions are theoretically in contravention of existing broad legal guidance on permissible network defence.⁴⁸

A government's position on the permissibility of hack back should be a function of two factors: permissible active defence techniques and the alleged identity of the adversary whose techniques are being mobilized against. Additionally, it is helpful to describe the permissibility of hack back in three broad categories: approved, permissible and forbidden:

- Active defence techniques span a gamut whose wide differentiation creates opportunity for policy consensus, from generally accepted techniques (like research on the tools and techniques of network intruders) to more invasive techniques where an organization is acting beyond the borders of its own networks. The use of extra-network techniques would normally be relatively infrequent vis-à-vis the measures an organization is empowered to implement within its own network. Furthermore, the prudent deployment of these techniques will typically require the active engagement of the highest levels of an organization's security, risk and legal leadership.
- It is helpful for policy-makers to establish clear guidance on the adversaries that attacked organizations are permitted to pursue. To underscore the point: hack back policy that does not require organizations to provide robust evidence to reliably establish the identity of the adversary they intend to pursue is inadvisable. It is important to recall that attributing an attack is difficult. But the inability to target a response, particularly if it acts beyond the borders of a network, risks creating enormous collateral damage. A perceived network intrusion could set off a cascade of reciprocal hack backs that may be destabilizing if the identity of the intruder is not well established.
- The alleged identity of the adversary should also affect the permissibility of active defence. Responding to a nation-state adversary may trigger significant collateral obligations for a host state of would-be active defenders. As such, policy-makers may consider curtailing attempts to attack nation-states. Policy-makers might also consider curtailing the use of active defence techniques against more sophisticated non-state adversaries, as those adversaries may have a greater ability to obfuscate their identity and dangerously escalate a conflict.


4.12 Active defence

Policy model: Active defence



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by active defence policy choices



Connecting policy to values

The value trade-offs created by active defence policy are shrouded in more ambiguity than other topics in cybersecurity. Little empirical evidence exists regarding the impact of active defence because it is not measured in most jurisdictions (owing to its questionable legal status).

- The starkest example illustrating the difficulty of understanding the trade-offs associated with active defence is its impact on security. In theory, the permissibility of more invasive active defence techniques should concern and deter adversaries as those adversaries will believe that the costs of criminal activity are higher. However, active defence might also reduce security for innocent bystanders, who may be the recipient of an incorrectly targeted hack back or, perhaps worse, experience collateral damage from an escalation of cyberattacks. In short, the use of more invasive active defence techniques has an ambiguous impact on security.
- The use of hack back techniques also has an ambiguous economic value in the long run. Even in the short term, the proliferation of such techniques will be costly as effective active defence is an expensive capability for an organization to field. In addition to the first order cost, active defence risks collateral damage, liability or an escalation of attacks in cyberspace.
- Enabling the private sector to act with a greater degree of freedom in cyberspace (embodied by a more permissive view of the active measures organizations may take to defend themselves) will increase the private sector's accountability to ensure security. With greater tools, they can rightly be expected to take a greater role in their own security. As active defence creates more private-sector accountability, it also creates substantial concerns for public-sector accountability. If an organization wrongfully responds to a nation-state, it is not clear what obligations the host state of the active defender has. In the case of a multinational, it might not even be understood which public sector is on the proverbial "hook". Is it the country of residence for the corporate headquarters or the country from where the attack was launched? Or is it both?
- To the extent that more organizations are empowered, active defence techniques are likely to be the province of very few organizations with significant capabilities (somewhat like attribution). Consequently, more permissible active defence policies are likely to decrease fairness.



4.13 Liability thresholds

Policy model

Private-sector organizations are increasingly subject to attacks of greater sophistication and persistence. The consequences of attacks are also becoming increasingly damaging; a digitally transformed business has more digital assets at risk. One increasingly difficult question confronting policy-makers is understanding how much risk should be borne by the public sector vs the private sector. Put differently: what is the reasonable duty of care that an organization should have? When does the public sector's obligation begin?

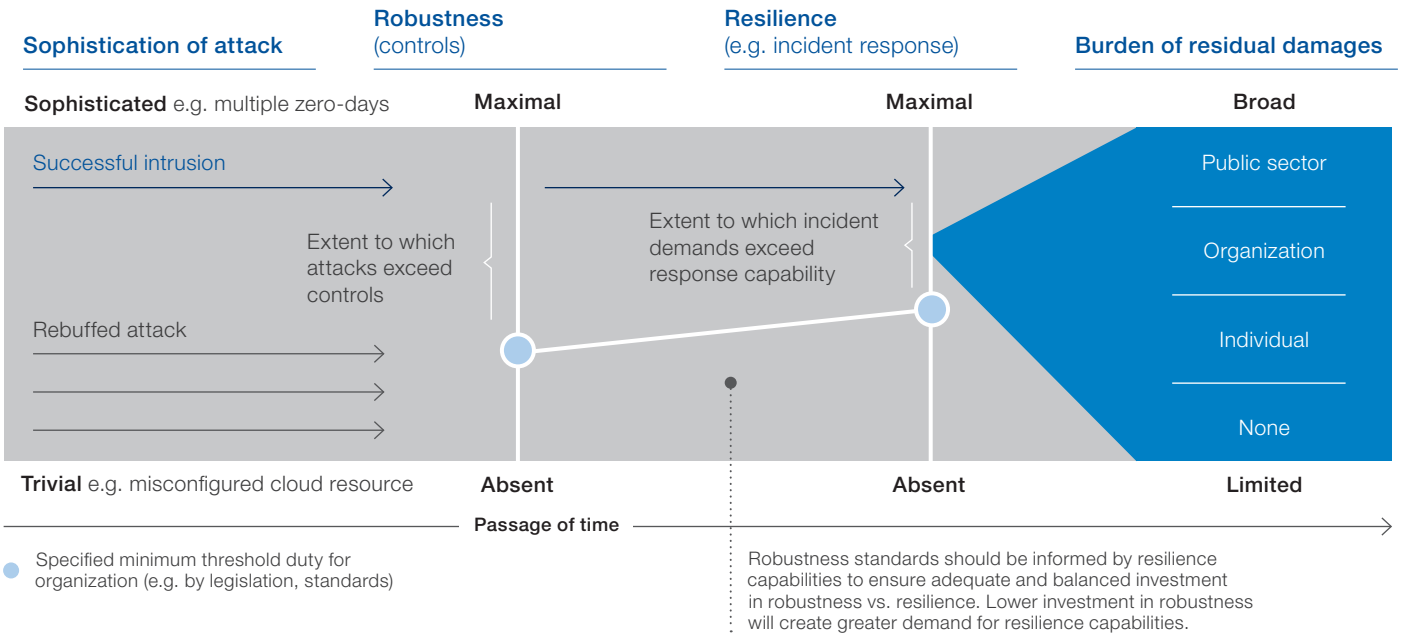
- The greater the duty of care an organization in the private sector needs to have, the more risk it needs to manage through investing in security technology, expertise, insurance (when possible) or adequate provisions (e.g. self-insurance). Given the increasing importance of insurance in managing risk, defining an organization's duty of care has consequences for cyber-risk bearing and the development of the adjacent insurance industry.
- A greater duty of care also has associated costs and, at some point, the incremental cost of additional security will fundamentally pervert an organization's business model. On the other hand, a more limited duty of care also has associated costs, potentially resulting in negligence that can be externalized.

No matter the duty of care an organization is expected to have, inevitably an attack will occur whose novelty and sophistication exceed established security controls, resulting in damages beyond the organization's prepared incident response abilities. Another key question is: what entity (if any) will bear the resultant residual damages occurring as a result of a successful intrusion?

- If those damages are borne by the targeted organizations, then the de facto impact will be to prompt that organization to review whether the resilience capabilities it has invested in are sufficient and, after quantifying its risk appetite, evaluate whether increased investment would be justified by diminished risks and costs.
- Alternatively, if those damages are borne by the public sector, downward pressure may be put on the resilience and robustness capabilities an organization develops. In practice, the fact that damages are being borne by the public sector does not necessarily imply that it will compensate those impacted by a successful incident, but may establish some minimum guarantee of protection to ensure trust (e.g. depositary insurance in the financial sector).

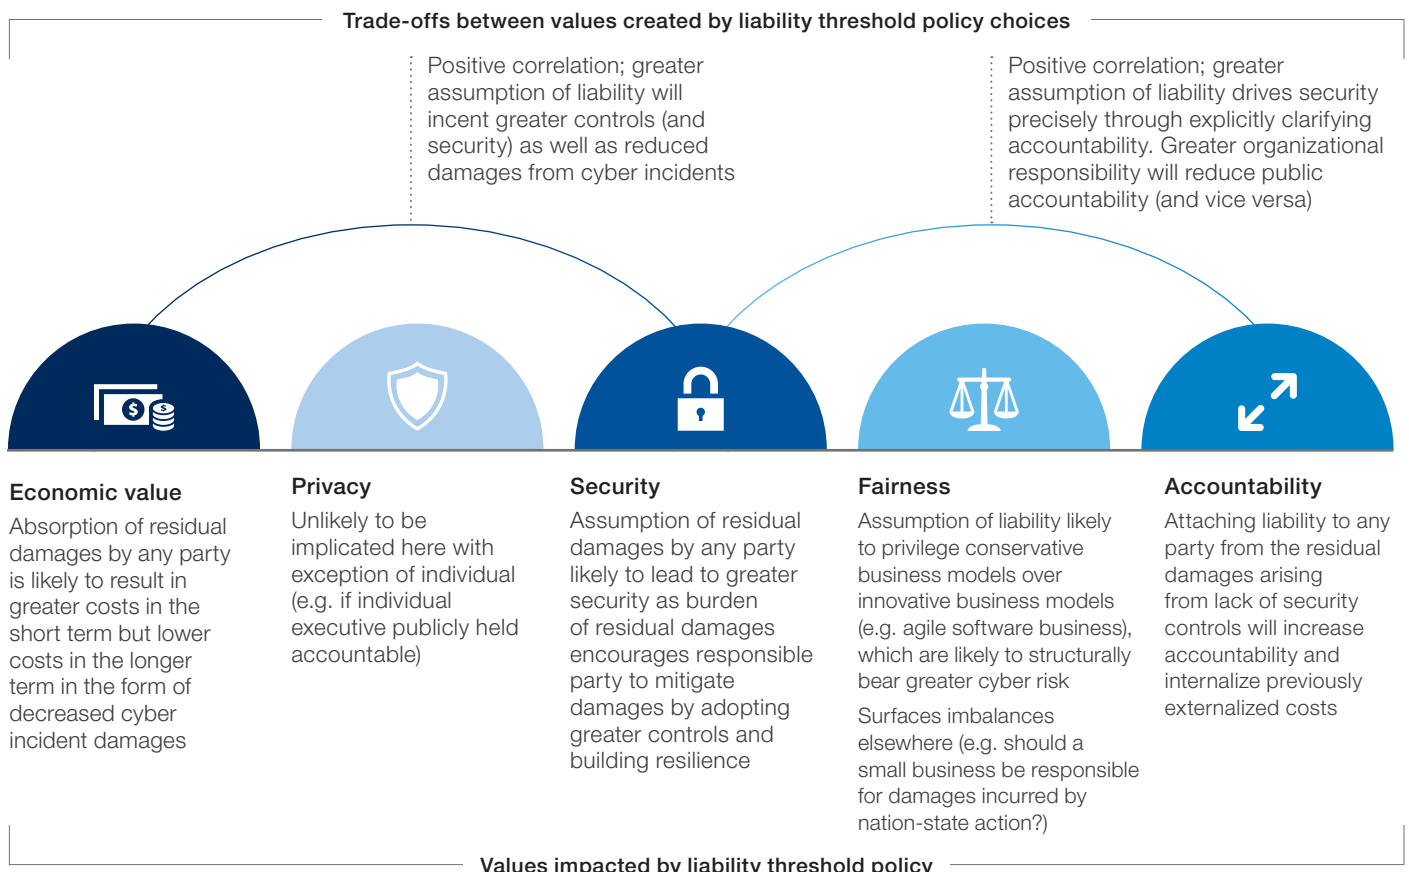
An important ancillary consideration is that no matter what duty is established, responsibility for assuring a predetermined level of robustness and resilience should be borne by the same entity. For example, for a given attack, if a business is wholly responsible for robustness and the public sector wholly responsible for resilience, then a business will be under-incentivized to invest in robustness. After all, it does not have to "clean up" the damage from an attack. Similarly, if an attack triggers the public sector's responsibility for robustness, the public sector should also be responsible for resilience (rather than distributing damages to businesses and their customers). Put simply, the entity empowered to act against a particular cyberattack should internalize the costs of its failure to thwart an attack.

Policy model: Liability thresholds



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by liability threshold policy choices



4.13 Liability thresholds

Case study: Publicly provided flood insurance

One analogue to separating robustness and resilience capabilities in cybersecurity is publicly subsidized flood insurance. In the United States, the government has extended subsidized insurance to homeowners in flood prone areas. Insurance is subsidized with the decidedly benevolent aim to help people recover from floods more quickly. However, one unintended consequence of that provision is that people are less incentivized to live outside of flood zones. After all, the costs of flooding are externalized from private individuals to the collective public. Consequently, damages caused by floods are higher than they would otherwise be.

Connecting policy to values

The core value most affected by the articulation of a duty of care is security:

- The greater the duty of care is expected of an organization, the more that organization will invest in both resilience and robustness capabilities. Furthermore, the assumption of residual damages from successful cyberattacks may drive improvements in security as organizations invest in measures to minimize those residual damages.
- To the extent the duty of care expected of organizations results in diminished cyberincident damages, it will also be economically beneficial. The benefits of a greater duty of care should become more apparent over longer periods of time as organizations work with insurers to develop a better understanding of how to manage a portfolio of cyber-risks through a combination of security controls and financial instruments.



4.14 Cyberinsurance

Definition

Cyberinsurance — a rapidly growing form of insurance for organizations seeking to manage cyber-related risks, such as first-party costs incurred as a consequence of a cyberattack, breach, business interruption, restoration and third-party liability; depending on the jurisdiction, regulatory fines/penalties may also be covered⁵⁰

Policy model

In addition to technical and behavioural measures, organizations are increasingly turning to cyberinsurance to help manage the financial consequences of cyber-related risk. Policy-makers are beginning to explore how cyberinsurance can not only help manage risk but incentivize mitigating it, as well. In an ideal world, insurers would offer cheaper insurance to companies contingent on better security controls. Insurers would also inform organizations seeking coverage about the controls they could implement to cost effectively reduce risk. So far, cyberinsurance is a nascent field and is offered by private companies, while policy-makers are experimenting with mandates and incentives to increase the adoption of insurance.

Policy-makers are confronted with two key questions pertaining to cyberinsurance (leaving aside the particulars of the coverage provided): what incentives, if any, should be offered to obtain insurance, and which entities should be prioritized for these incentives?

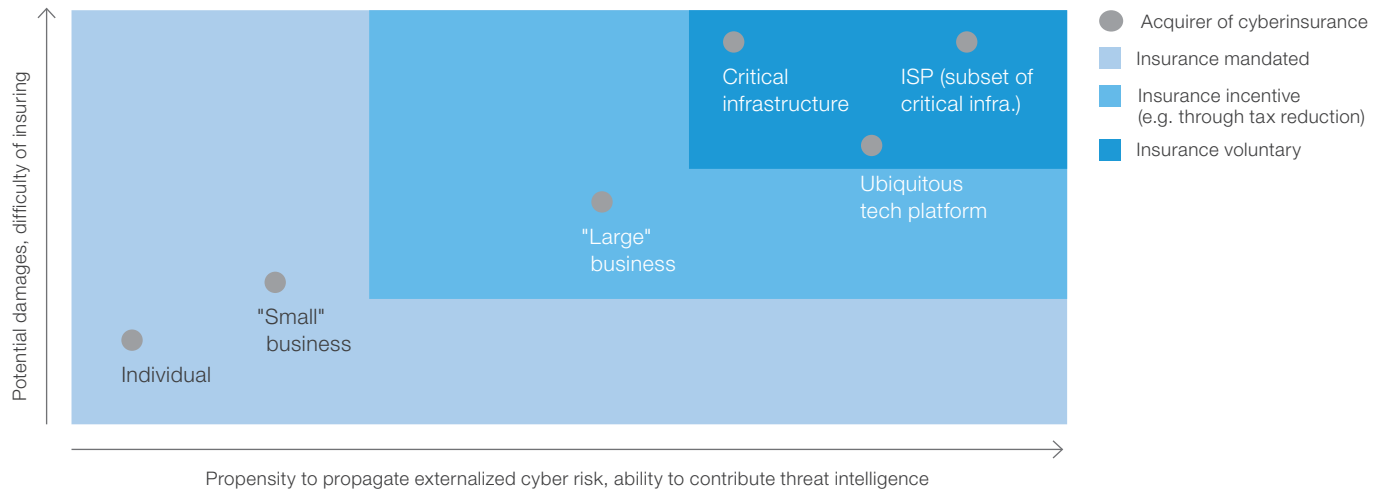
- Broadly speaking, a state may intervene at three levels in the insurance market, with increasing likelihood of private-sector adoption but increasing costs, as well: voluntary (no incentives); incentivized (e.g. tax deduction); and mandated insurance. If no incentives for insurance exist, the upfront costs are likely to be low but, in the long run, depending on how liability is defined, at some point cyber costs will be borne in an outsized fashion by some entity either in the private or public sector. These costs are likely to be greater in the absence of the security control adoption promoted by insurance. On the other end of the spectrum, an insurance mandate will lead to greater upfront costs for the private sector but to smaller costs in the long run as companies adopt security controls to minimize insurance costs.
- A number of entities could be targeted for state-incentivized insurance. Given finite resources, it may be more valuable to target insurance incentives towards organizations that are less mature and capable of weathering the financial consequences of cyberattacks.

The provision of cyberinsurance is not an unalloyed collective good even if insurers incentivize adequate cybersecurity risk mitigation. The insurance industry itself must be carefully monitored for the buildup of financial risk associated with bearing the costs of cyberincidents. In some jurisdictions, regulators have been concerned by the size of the implicit liability borne by insurers underwriting cyber-risks (also known as “silent” risk).⁵¹

4.14 Cyberinsurance

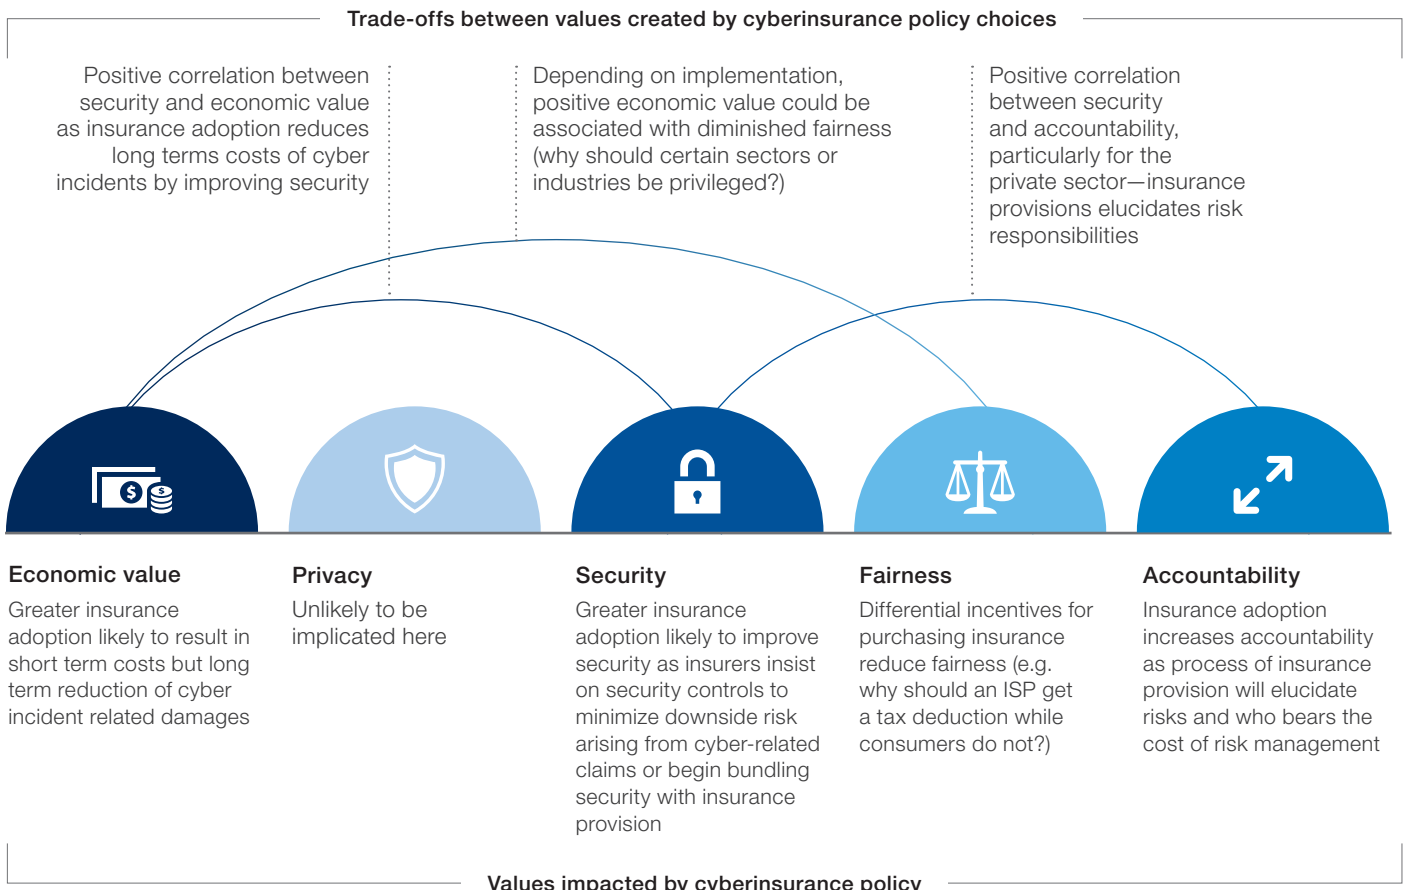
Policy model: Cyberinsurance

Prioritizing who should acquire cyberinsurance



Policy model described herein is illustrative of one of many different viable policy configurations.

Key values trade-offs created by cyberinsurance policy choices



Case study: Department of Homeland Security, Cyber Incident Data and Analysis Repository (CIDAR)

The Cyber Incident Data and Analysis Repository (CIDAR) is an initiative led by the U.S. Department of Homeland Security to solve one of the key constraints to cyberinsurance adoption: data — in particular, data that connects the failure of a specific security control with the damages incurred as a consequence. Without data, insurers cannot price the risk a given organization presents and thus cannot offer insurance in a robust way (i.e. they must offer it at such a steep price that few organizations can afford to adopt it or, perhaps worse, the few that purchase the insurance are the equivalent of cybersecurity “lemons”).⁵²

CIDAR helps illustrate how very important and vexing data limitations are for the maturation of the cyberinsurance market. Relative to other types of risks insurers’ cover, cyber-risk is very difficult to measure and price. It is difficult to measure for three reasons. First, there is limited historical data. Insurers do not have a reliable indicator of the damages associated with the failure of security controls. Second, cyber-risk is “fat tailed”; very extreme events tend to occur somewhat more commonly than one would expect. Statisticians have difficulty measuring fat-tailed risk and thus, relative to other risk, cyber-risk requires a relatively larger sample size to confidently assess. Third, and perhaps most frustratingly for insurers, cyber-risk measurements are subject to an inherent uncertainty associated with threat vectors changing over time. Damage estimates associated with the failure of corporate PC-centric security controls in the early 2000s were unlikely to be adequate for assessing a bring-your-own-device corporate environment in 2012, and are even less meaningful for assessing a workplace blanketed with smart sensors in 2017. Cyberinsurance provision is hindered by the fundamental paradox of peering backward at an incomplete history to estimate forward-looking future technology risks.

Connecting policy to values:

It is difficult to predict the normative trade-offs that will result as a consequence of policy choices impacting cyberinsurance, given the industry’s relative nascence. But in general, policies that promote increased adoption of cyberinsurance should lead to improved security as companies gain a better understanding of their own cybersecurity risk profile. The more data insurers have, the better they should be able to assess the relative importance of different risks, and price insurance accordingly.

Risk transparency also helps promote greater private-sector accountability. An organization aware of how it can act to mitigate its own risks should be held to a higher standard.

Over time, increased insurance adoption should lead to decreased security-related costs (inclusive of insurance), given the ability to reduce a given company’s risk profile.

5. The future of cyber resilience

The frameworks and discussions outlined in this document endeavour to provide the basis for fruitful collaboration between the public and private sectors in securing shared digital spaces. In the coming years, the World Economic Forum will continue to offer insights and spur action in this area as cyber resilience remains a top-of-mind topic for decision-makers. The aim is to further efforts to advance cybersecurity.

Appendix: Normative trade-offs framework

Developed by the System Initiative on Shaping the Future of Digital Economy and Society as a tool for removing the veil of ambiguity from difficult decisions

In a number of contexts, from business to politics to the social sector, leaders have to make decisions and prioritize one set of values over another — a policy-maker may be forced to choose between allocating a national budget towards education versus healthcare; a business leader may be forced to choose between capturing market share versus profitability.

To make these decisions, leaders often seek out data to inform their choices. For example, policy-makers have reams of budget analyses and business leaders have granular visibility into customer segments. While this decision-making environment is rich in facts that may support any decision, the process itself is often divorced from the core values leaders are attempting to promote and prioritize.

However, these hard-to-quantify values often implicitly frame the terms of the debate through which policy is made. To facilitate informed decision-making on these “soft” questions, a three-part decision framework was developed, which has been implemented across numerous efforts in the World Economic Forum System Initiative on Shaping the Future of Digital Economy and Society, including cyber resilience, and which is intended for broad dissemination.

The objective of this framework is to surface the values that underlie decision-making and offer a transparent and collaborative process by which leaders can make and explain policy decisions with normative implications. This framework is comprised of three steps:

1. Articulate the option space
2. Isolate the most important values
3. Quantitatively rank feasible choices

1. Articulate the option space

To make well-informed decisions where determinations of value trade-offs are required, it is necessary to have a firm grasp of the full set of options and key elements that help distinguish one possible decision from its “nearest neighbour”.

A full set of options, unconstrained by the limitations of present circumstance, helps push the boundaries of thinking. At this stage, it is important to consider many “possible” decisions in a policy space, even if some may be undesirable or implausible. This exercise allows for later attribution of values or norms to be clearer and more explicit.

In cybersecurity policy-making, one commonly debated issue is the handling of so-called “zero-days”. Zero-days are exploitable vulnerabilities not known about publicly (they are in “day 0” of their discoverability). These vulnerabilities (and exploits which take advantage of them) can be catalogued and stockpiled by national defence organizations and deployed offensively. These zero-days can also be shared with the software vendors whose product is vulnerable, so they can develop measures to mitigate and patch these vulnerabilities.

In the process of elucidating the full option-space for zero-days, the Working Group convened by the World Economic Forum suggested that focusing on the government’s role in developing and sharing zero-days, while important, is a reactive and limited policy posture. After all, they reasoned, a software vulnerability first has to be coded before a debate can arise about how to share knowledge of that vulnerability to promote competing valid national interests. In brief, in articulating the full set of areas where policy-makers could contribute, it became obvious that much of the debate — while valid — did not adequately consider other important elements.

Appendix: Normative trade-offs framework

2. Isolate the most important values

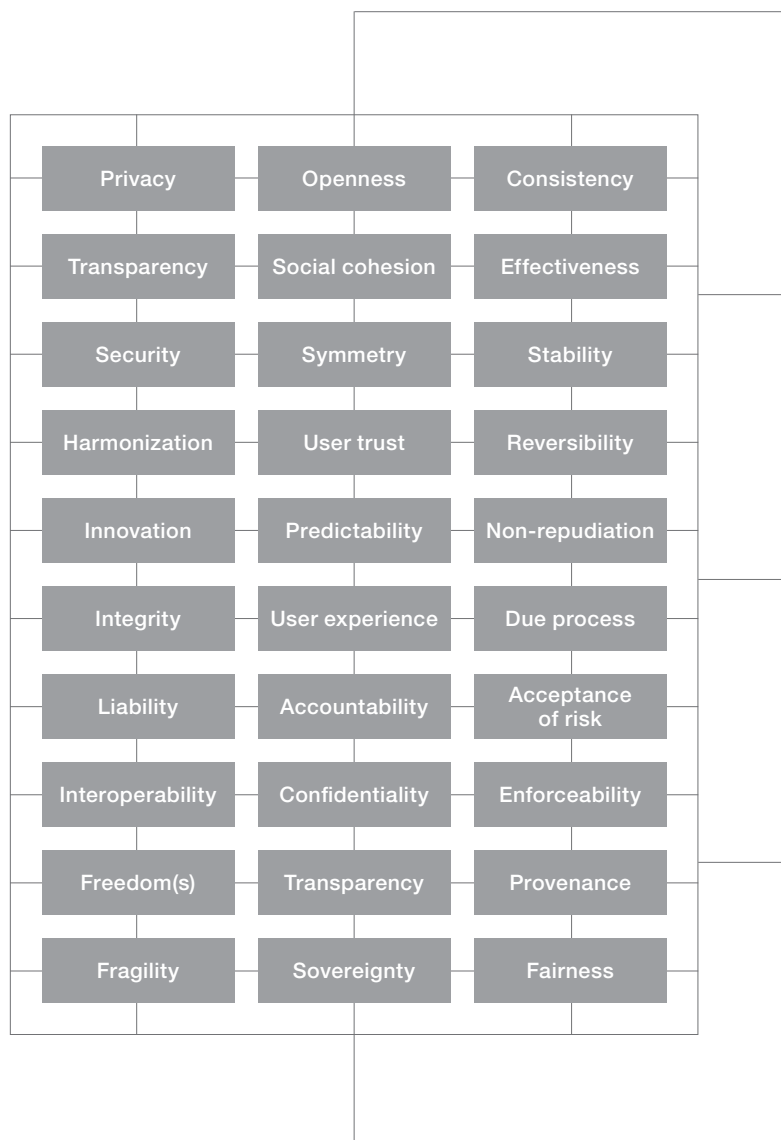
After mapping the option-space, it is necessary to develop the “long list” of values to consider in the process of making a decision. It is meant to be a list of all the values that might be held by a given constituency with respect to a policy area. The list is not meant to be exhaustive, but should include a sufficient number of values to ensure that the most important or most likely to give rise to a values conflict are represented. Depending on the context, care should be taken to ensure that the values described are relevant to the various political, cultural and personal differences among stakeholders liable to be affected by the decisions in question.

For the Playbook for Public-Private Collaboration, the Forum convened a group to outline the key values that policy-makers should weigh in making choices between different cybersecurity policy options.

After defining the “long list of values”, the Working Group began simplifying and aggregating these values to a tractable and complete set. Again, taking the example of a recent discussion on cybersecurity policy, the more than 20 values that were initially identified as significant were eventually pared down to a list of five key values animating policy debate: security, privacy, fairness, economic value and accountability.

- **Security** — the protection of assets (tangible and intangible) from damage. Assets may be anything of value, including the well-being of individuals. Damage may comprise the loss of availability, integrity and, where applicable, confidentiality of assets resulting in a diminution of value for the rightful owners of the asset.
- **Privacy** — the ability of an individual, group or organization (e.g. business) to limit information about themselves. The boundaries of privacy vary by context and by country. The domain of privacy partially overlaps with security (confidentiality), which can include the notion of appropriate use as well as protecting information.
- **Fairness** — the extent to which entities within a given nation-state will be impacted symmetrically (or with otherwise perceived appropriateness) by policy, including due process. Perceptions of appropriateness will vary by context and by country.

Beginning consideration of values for cybersecurity policy




Concluding consideration of values for cybersecurity policy



Economic value

Effectiveness
Innovation
Interoperability
Sovereignty
Social cohesion
User experience



Privacy

Acceptance of risk
Freedom(s)
Confidentiality



Security

Consumer trust
Reversibility
Integrity
Availability
Non-repudiation



Fairness

Symmetry
Due process
Openness



Accountability

Harmonization
Acceptance of risk
Predictability
Stability
Provenance
Consistency
Transparency

- **Economic value** — the amount of monetary and common wealth, and commerce statically (e.g. current market participants) and dynamically (e.g. in the future from innovation) resulting from, or destroyed by, a given policy choice. Lower costs from cyberincidents may also contribute to greater economic value.
- **Accountability** — the extent to which an entity (individual, group, organization) can be held responsible or even liable for consequences arising out of its action or inaction. Public- and private-sector accountability have been separately delineated to demonstrate how burden shifts in particular policy models.

In addition to analytical tractability, the forcing function of shaping values is itself informative about how to think about value-based decisions:

- Not all values are equally relevant or important for a given policy discussion. For example, security is qualitatively more important as a dimension to evaluate cybersecurity policy than interoperability.
- Some values subsume others in their scope. For example, innovation is a subset of economic value.
- Some values enable others but are not fundamentally important in themselves. For example, transparency has little intrinsic importance but is enormously empowering to greater accountability.

3. Quantitatively rank feasible choices

After defining the policy and business choices a leader can make on a given topic and the values that should be considered in making those decisions, one can begin confining the option-space; certain choices simply cannot be made by virtue of a fundamental constraint. For example, while it is important to consider a world in which many entities can fully monitor internet traffic, in practice the cost of capturing and effectively analysing such a massive volume of data will be prohibitive for most governments.

Having pared down the option-space into a set of feasible choices, it is important to explicitly enumerate the risks and benefits associated with a given choice.

Appendix 1: Normative trade-offs framework

Next, to ensure that subjective values are thoroughly debated and understood, it is valuable to numerically rank how much each value is promoted. Assigning a numerical estimate to how much a value is promoted or prioritized serves another important forcing function. By assigning a number to a given value, organizations are forced to make a more granular and nuanced judgement as to the impact of a given choice. Such quantification (even if only for illustrative purposes) also avoids absolutist justifications of preferred policy options and false binaries.

For example, in the context of cybersecurity and the values that different policy choices embody, a persistent problem is stakeholders grasping for rhetorical simplicity. For example, defence ministries will often argue that absent security, no other liberties can be secured. But the rhetorical simplicity of such an argument is undercut by being forced to articulate numerically the relative difference of different policies on security. If a policy is indeed able to provide significantly enhanced security, it should be easy to articulate either through anecdotal evidence or, better yet, numerical evidence.

The choice of numerical ranking is also important. A numerical scale with too many degrees of gradation will be intellectually taxing. Choosing a numerical scale that is odd numbered (e.g. with five options) risks allowing clustering of evaluations to form around the number 3. And a lukewarm indicator of a given choice's impact is less valuable (e.g. 3 in the context of a 1 to 5 scale where 1 is the lowest prioritization of a given value and 5 is the highest prioritization of a given value). Just as exploitable differentiation is key to statistical inference, differentiation draws into high relief the trade-offs decisions require.

Another important benefit of forcing a numerical thinking for decision-making is its ability to illuminate inconsistencies or themes across different questions that a leader in a given organization will confront. For example, in the course of defining the numerical impact of policy choices, the World Economic Forum cyber resilience project found that the normative impact of insisting on weak encryption for companies in the private sector is similar to the normative impact of allowing employers to monitor the internet traffic of their employees. For most participants, the intellectual resemblance between these policies was not evident until this exercise was completed.

In the end, this exercise can be distilled to a series of “if ..., then ...” statements of the type “if a decision-maker prioritizes x value, then he/she should most likely promote y policy option.” These statements form the basis for a values-focused set of decisions and for a rubric to measure current policy decisions vis-à-vis professed values.

When to use a decision framework on values

A decision framework for normative questions is useful — it helps force relevant conversations quickly and, in imposing rigour on a typically circuitous process, helps ensure that there is forward movement on the outcome: making a decision.

However, the use of a decision framework implicitly prioritizes deliberation. Discussions of values are cognitively taxing and take time. In some contexts, the ability to rapidly make a decision may obviate the need for a well-considered framework, particularly if those decisions are easily reversible.

Acknowledgements

The World Economic Forum cyber resilience project is a global, multi-industry, multistakeholder endeavour aimed at contributing to a safer and stronger connected society by improving cyber resilience. The project engages stakeholders across multiple industries and governments from around the world.

The governance and strategic direction for this project is provided by the System Initiative on Shaping the Future of Digital Economy and Society and our dedicated Steering Committee and Working Group. This Playbook is based on numerous discussions, workshops and research. The aggregated opinions expressed herein may not necessarily correspond with each and every view stated in the project. The project intended to seek consensus among those participating on different areas of expertise.

Sincere thanks are extended to the experts who contributed their unique insights to this Playbook. On this, grateful appreciation for its generous commitment and support goes to The Boston Consulting Group.

Steering Committee

The Steering Committee brought together senior and experienced leaders with a variety of backgrounds to provide broader strategic guidance on the topics addressed. Steering Committee members also

The World Economic Forum looks forward to these efforts continuing at the Global Centre for Cybersecurity under the leadership of Alois Zwinggi, Member of the Managing Board, and Ushang Damachi and Karen Wong. This project benefited greatly from the support and commitment of the Technology, Media and Digital Industry Community, led by Alan Marcus, and the ICT Industry Community, led by Danil Kerimi, Lauren Joseph and Adam Sherman. Thanks also go to Adam Schlosser, Digital Trade and Data Flows Project Lead, for his contributions to the section on cross-border data flows. Finally, this project would not be possible without the leadership of Derek O'Halloran, Head of the System Initiative on Shaping the Future of Digital Economy and Society, as well as the support of the Digital Economy and Society team — Manju George, Kelly Ommundsen and Justine Moscatello — and the invaluable input and guidance provided by the Forum's Managing Board and Chairman.

This Playbook is dedicated to the memory of Jean-Luc Vez, Head of the Global Centre for Cybersecurity. Jean-Luc will be remembered for his deep commitment to partnerships for improving global security, his leadership in our cyber efforts, and his kindness and warmth.

Daniel Dobrygowski

Lead, Cyber Resilience project

dedicated significant time and effort, volunteering their own expertise, insight and deep relationships within their respective spheres.

Elizabeth Joyce	Vice-President and Chief Information Security Officer	Hewlett Packard Enterprise	USA
David Koh	Chief Executive	Cyber Security Agency	Singapore
Cheri McGuire	Group Chief Information Security Officer	Standard Chartered Bank	United Kingdom
Bradford L. Smith	President and Chief Legal Officer	Microsoft	USA
Amy Weaver	President and General Counsel	Salesforce	USA

Acknowledgements

Expert Working Group

The Expert Working Group brought together leading academic experts, thinkers and senior executives from across industries and sectors. Working Group members dedicated their time and other resources to help

the project cover the range of functional perspectives that needed to be integrated in this topic, in particular risk, security, technology and legal perspectives.

Bénédicte Suzan	R&T and Innovation Management Airbus Defence and Space	Airbus	France
David O'Brien	Senior Researcher	Berkman Klein Center for Internet & Society	USA
Phillip Harrington	Senior Managing Director	Brock Capital Group	USA
Michael Nelson	Public Policy	Cloudflare	USA
Sithuraj Ponraj	Deputy Director, International Cooperation & Partnership Office	Cyber Security Agency	Singapore
Gadi Evron	Chief Executive Officer	Cymmetria	Israel
Andy Radle	Chief Architect for Cloud Security	Hewlett Packard Enterprise	USA
Nat Mokry	Senior Director of IAM & Industry Next Security	Hewlett Packard Enterprise	USA
George D. DeCesare	Senior Vice-President, Chief Technology Risk Officer	Kaiser Permanente	USA
Anton Shingarev	Vice-President, Public Affairs	Kaspersky Lab	Russia
Paul Nicholas	Senior Director, Global Security Strategy and Diplomacy	Microsoft	USA
Lindsey Finch	Senior Vice President & Associate General Counsel, Global Privacy & Product	Salesforce	USA
Marc Porret	ICT Coordinator	United Nations CTED	USA
Michael Nunes	Innovation and Technology Policy	Visa	USA
Lori Bailey	Global Head of Cyber Risk	Zurich Insurance Group	USA

The project team would also like to thank the senior leaders who supported its work by providing input, expertise and thoughtful commentary during the project development: Keith Alexander (USA), Edward Amoroso (USA), William H. Saito (Japan), Toshikazu

Okuya (Japan), Bjarne Eckardt (Europe), Jeff Prall (USA), Philippe Cotellet (France), Sadie Creese (UK), Yosi Shneck (Israel), Jim Pinter (USA), Haizhou Gu (UN) and Han Soal Park (UN).

Project Strategy Adviser: The Boston Consulting Group

David Mkrtchian	Consultant (Seconded to the Forum)	USA
Walter Bohmayr	Senior Partner and Managing Director	Austria
Stefan Deutscher	Associate Director	Germany

Special thanks to the following Boston Consulting Group experts for their insightful contributions to the cyber resilience effort and this report: Mark Connelly, Chief Information Security Officer; Michael Coden, Managing Director, BCG Platinion; Nadya Bartol, Associate Director BCG Platinion; Astrid Blumstengel,

Global Market Director, Technology Advantage; Shoab Yousuf, Principal; Troy Thomas, Associate Director; Gregory Boison, Associate Director; Alexander Tuerk, Project Leader; Andrew Smolenski, Project Leader; Alex Asen, Global Knowledge Team Lead, Cybersecurity; Jin Bo, Senior Knowledge Analyst.

Contacts

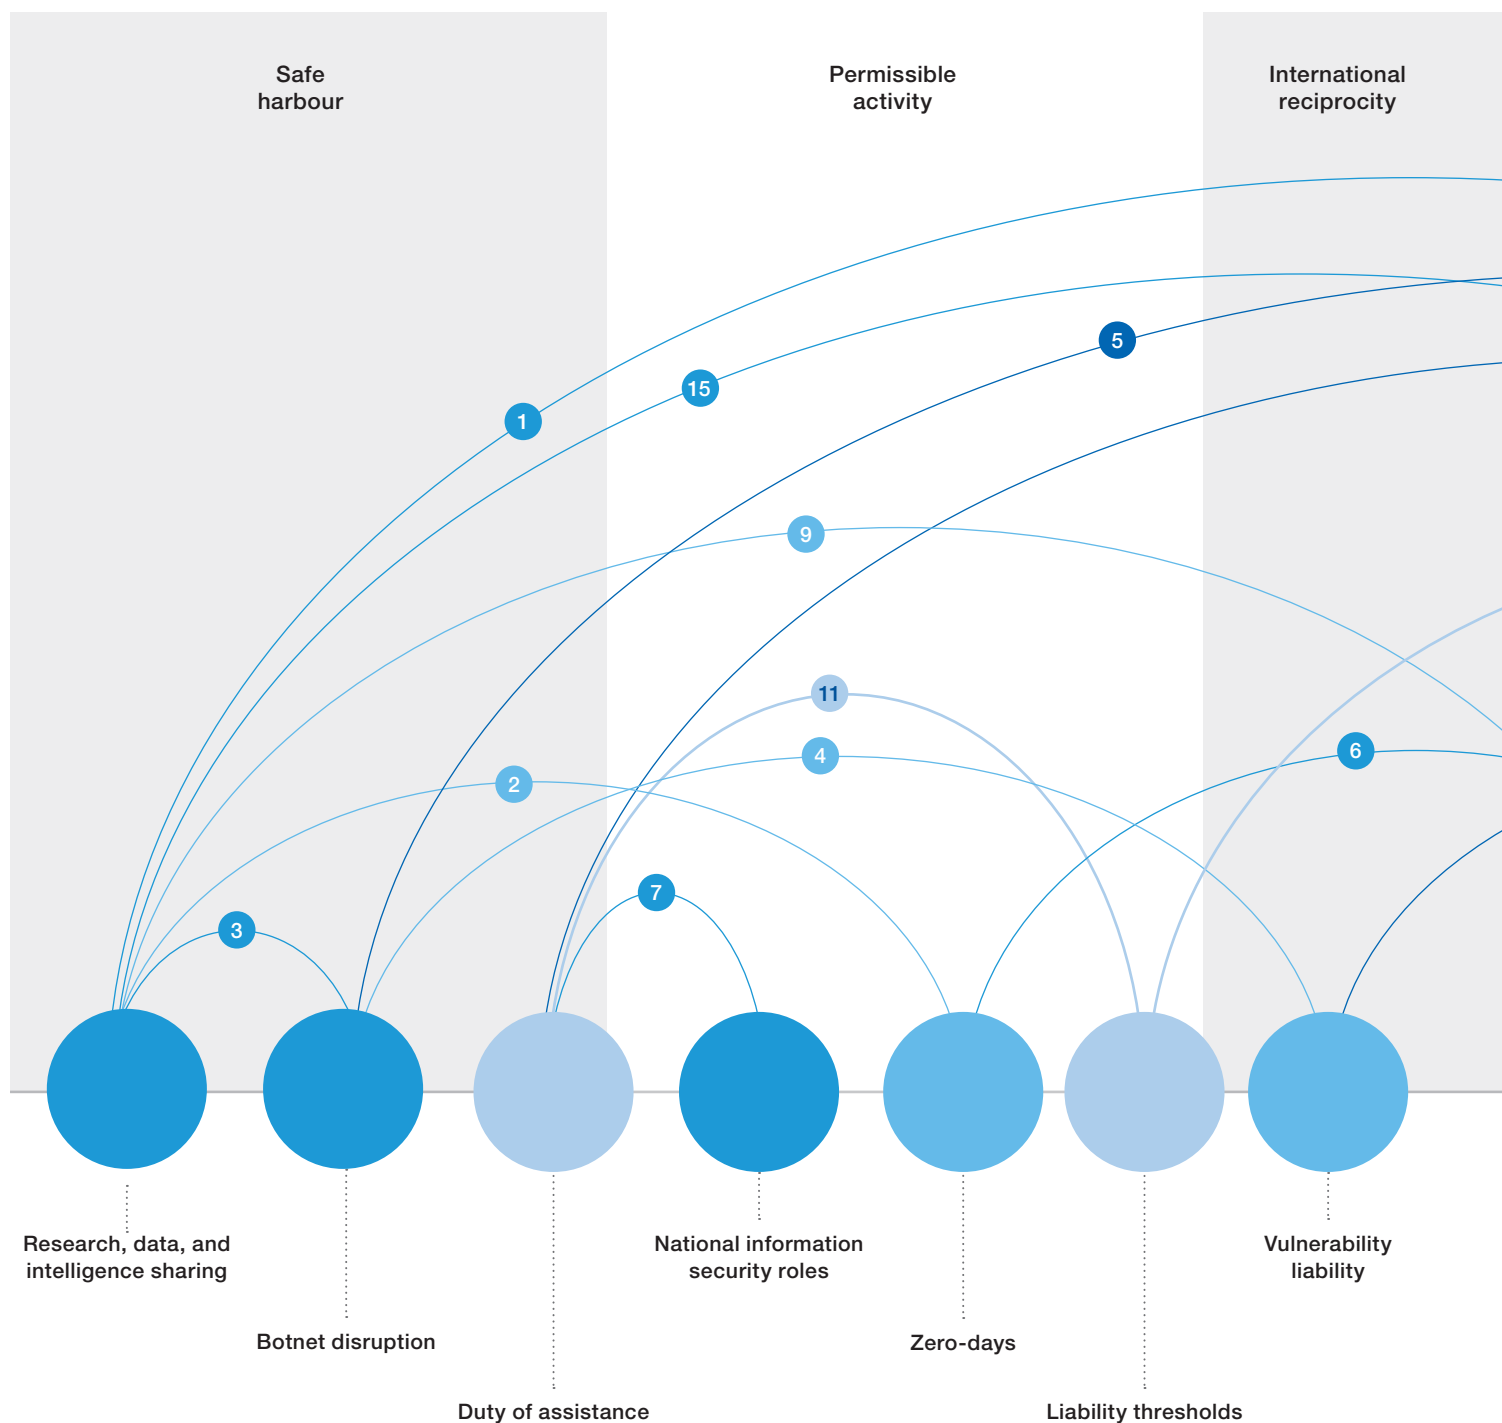
Derek O’Halloran

Head of the System Initiative on Shaping the Future of Digital Economy and Society
World Economic Forum

Daniel Dobrygowski

Global Leadership Fellow, Project Lead
World Economic Forum

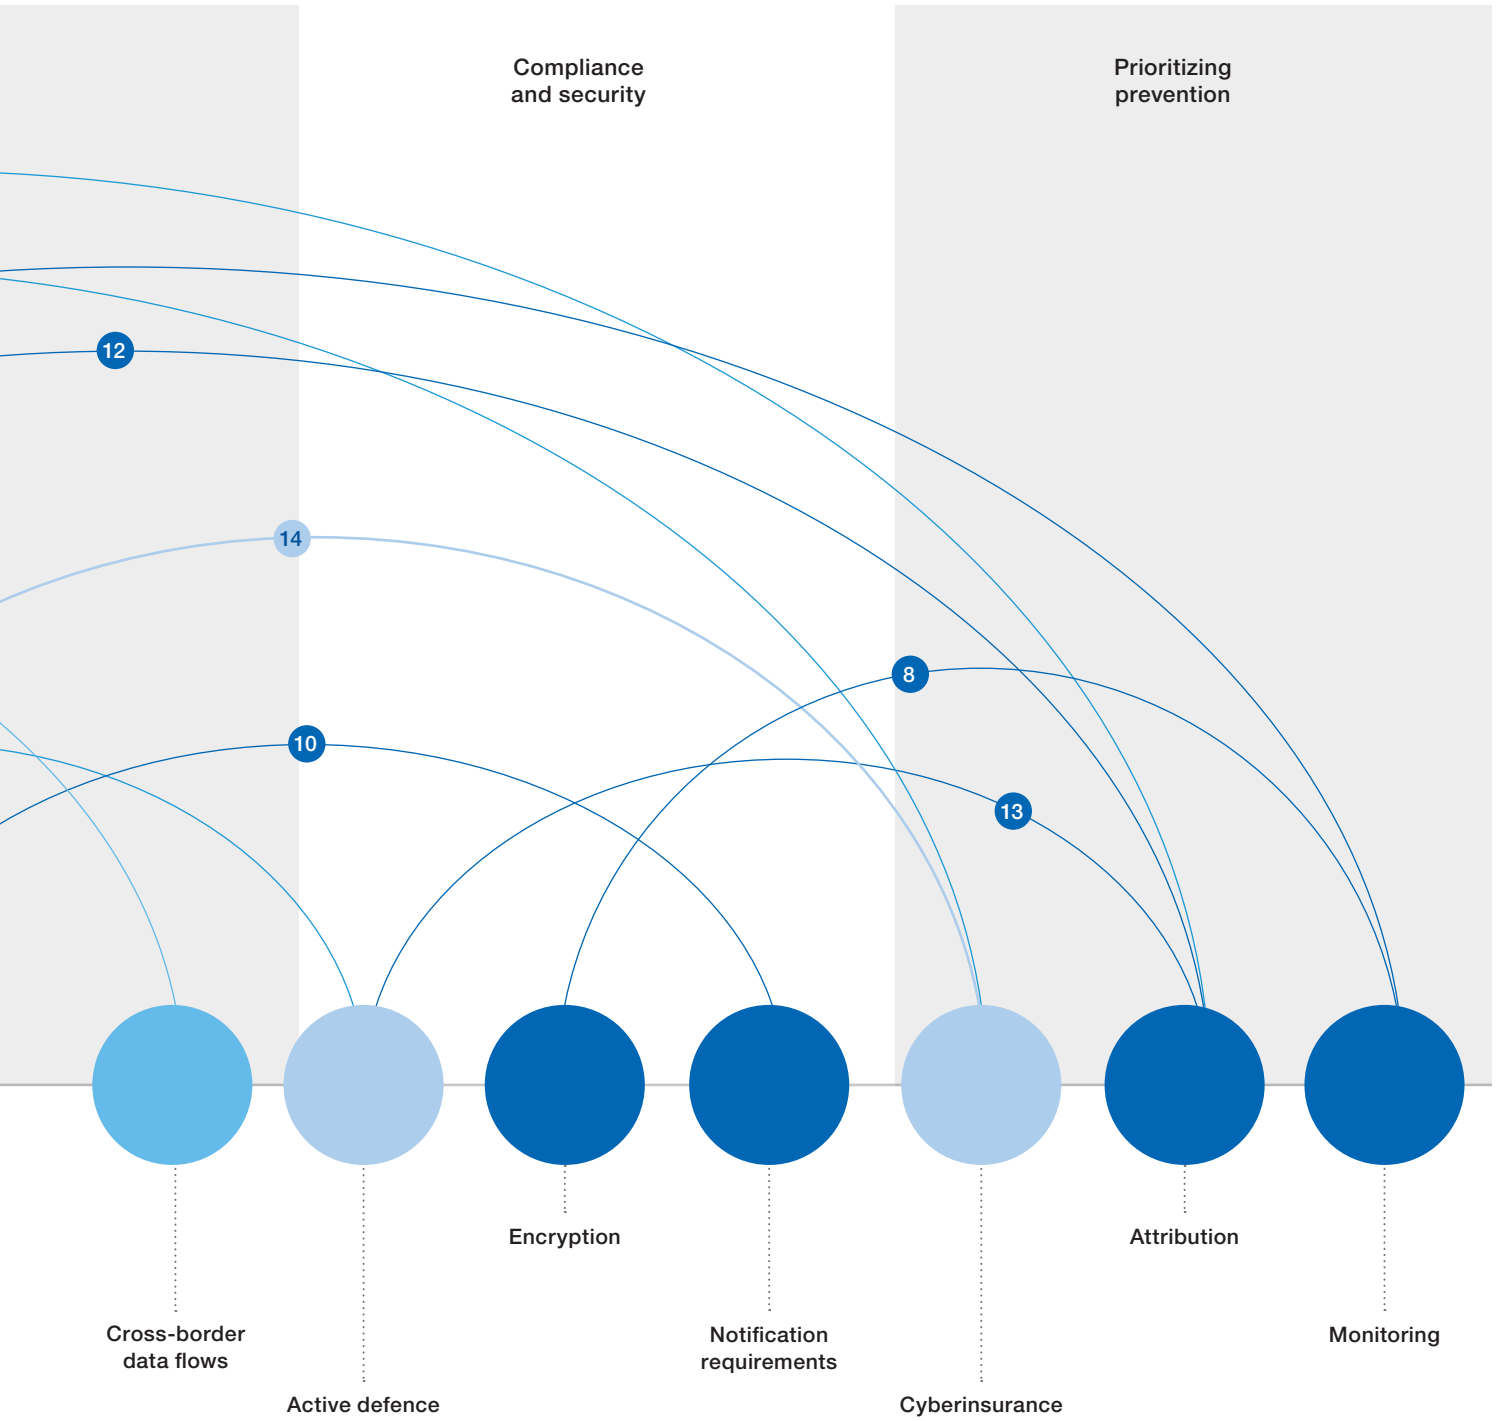
CyberResilience@weforum.org



Key linkages between policy topics

- 1 Attribution key element of Intelligence, particularly for public sector
- 2 Zero-day vulnerabilities crucial opportunity for governments to share threat intelligence
- 3 Botnet disruption facilitated by rapid and well-coordinated research and action
- 4 Securing vulnerabilities through avoidance or patching may diminish threat surface for botnet operators
- 5 More invasive monitoring capabilities may allow ISPs to police botnet more effectively
- 6 Extent of active defence permitted by private sector key element of national roles and responsibilities
- 7 Granular understanding of government duty of assistance fundamental to national cyber resilience
- 8 Greater adoption of strong encryption will hinder the ability to monitor network traffic

Note: List of connections between topics not exhaustive.



- 9 Limitations on cross-border data flows may introduce friction into intelligence sharing
- 10 Heightened notification requirements may result in increasing investment to secure known vulnerabilities
- 11 Duty to assist integrally linked with liability—where private sector cannot be reasonably expected to secure, government steps in
- 12 Nation-state attribution may trigger government duty to assist the private sector

- 13 Active defence may result in collateral damage without well-defined attribution and safeguards (e.g. organization vs. nation-state)
- 14 Liability thresholds circumscribe the nature of cyberinsurance incentivized
- 15 Cyberinsurance can be more effectively priced and deployed given greater data and intelligence

Endnotes

- ¹ Electronic Frontier Foundation. Barlow, J. P. (2017, 16 May). "A Declaration of the Independence of Cyberspace". Retrieved 11 December 2017 from <https://www.eff.org/cyberspace-independence>
- ² Wired. Gates, B. (2002, 17 January). "Bill Gates: Trustworthy Computing". Retrieved 11 December 2017 from <https://www.wired.com/2002/01/bill-gates-trustworthy-computing/>
- ³ Ministry of Foreign Affairs of the People's Republic of China. "Remarks by H.E. Xi Jinping, President of the People's Republic of China At the Opening Ceremony of the Second World Internet Conference", 16 December 2015, http://www.fmprc.gov.cn/mfa_eng/wjdt_665385/zyjh_665391/t1327570.shtml
- ⁴ TechCrunch. Biggs, J. (2017, 25 July). "Hungarian hacker arrested for pressing F12". Retrieved 11 December 2017 from <https://techcrunch.com/2017/07/25/hungarian-hacker-arrested-for-pressing-f12/>
- ⁵ Zerodium, "The premium acquisition program for zero-day exploits and advanced cybersecurity research". (n.d.). Retrieved 19 December 2017 from <https://zerodium.com/>
- ⁶ Harvard Kennedy School Belfer Center for Science and International Affairs. (n.d.). "Taking Stock: Estimating Vulnerability Rediscovery". Retrieved 11 December 2017 from <https://www.belfercenter.org/publication/taking-stock-estimating-vulnerability-rediscovery>
- ⁷ Open Web Application Security Project (OWASP). (n.d.). "Welcome to OWASP". Retrieved 12 December 2017 from https://www.owasp.org/index.php/Main_Page
- ⁸ National Institute of Standards and Technology (NIST). Kissel, R. et al. (October 2008). Security Considerations in the System Development Life Cycle. Gaithersburg, MD: U.S. Dept. of Commerce, Technology Administration, NIST
- ⁹ Slate. Gallagher, R. (2013, 16 January). "Cyberwar's Gray Market: Should the secretive hacker zero-day exploit market be regulated?" Retrieved 19 December 2017 from http://www.slate.com/articles/technology/future_tense/2013/01/zero_day_exploits_should_the_hacker_gray_market_be_regulated.html
- ¹⁰ Fortune. (n.d.). "Google's Elite Hacker SWAT Team vs. Everyone". Retrieved 11 December 2017 from <http://fortune.com/2017/06/23/google-project-zero-hacker-swat-team/>
- ¹¹ Symantec Corporation. (February 2011). Falliere, N. et al. "W32.Stuxnet Dossier." Retrieved 19 December 2017 from www.symantec.com/content/en/us/enterprise/media/security_response/whitepapers/w32_stuxnet_dossier.pdf
- ¹² iPhoneDevSDK. (October 2008). "Average time spent creating an app, poll". Retrieved 12 December 2017 from <http://iphonedevsdk.com/forum/iphone-sdk-development/3948-average-time-spent-creating-an-app-poll.html>
- ¹³ Inside EVs. "Infographic: Chevy Volt Has 10 Million Lines of Code; F-22 Raptor Only Has 1.7 Million". (n.d.). Retrieved 12 December 2017 from <http://insideevs.com/infographic-chevy-volt-has-10-million-lines-of-code-f-22-raptor-only-has-1-7-million/>
- ¹⁴ Slashdot. (2013, 1 February). "Mars Rover Curiosity: Less Brainpower Than Apple's iPhone 5". Retrieved 12 December 2017 from <http://slashdot.org/topic/bi/mars-rover-curiosity-less-brainpower-than-apples-iphone-5/>
- ¹⁵ Open Hub, Black Duck Software, Inc. (n.d.). "Chromium (Google Chrome)". Retrieved 19 December 2017 from www.openhub.net/p/chrome. Code count as of 2013 (6.7M)
- ¹⁶ Clarke, G. (2011, 22 September). "CERN's boson hunters tackle big data bug infestation". Retrieved 12 December 2017 from http://www.theregister.co.uk/2011/09/22/cern_coverity/
- ¹⁷ Wired. Newcomb, D. (2012, 3 December). "The Next Big OS War Is in Your Dashboard". Retrieved 12 December 2017 from <http://www.wired.com/autopia/2012/12/automotive-os-war/>
- ¹⁸ Microsoft, Microsoft On the Issues. (2013, 29 October). "New cybersecurity report details risk of running unsupported software". Retrieved 19 December 2017 from <https://blogs.microsoft.com/on-the-issues/2013/10/29/new-cybersecurity-report-details-risk-of-running-unsupported-software/>
- ¹⁹ SecurityScorecard. (2016, 8 June). "How Big is the End-Of-Life Cybersecurity Problem?" Retrieved 19 December 2017 from <https://securityscorecard.com/blog/end-of-life-cybersecurity-infographic>
- ²⁰ The Hacker News. Kumar, M. (2014, 14 April). "HeartBleed Bug Explained - 10 Most Frequently Asked Questions". Retrieved 12 December 2017 from <https://thehackernews.com/2014/04/heartbleed-bug-explained-10-most.html>
- ²¹ Recorded Future. (2017, 19 October). "The Dragon Is Winning: U.S. Lags Behind Chinese Vulnerability Reporting". Retrieved 12 December 2017 from <https://www.recordedfuture.com/chinese-vulnerability-reporting/>
- ²² TechCrunch. (2016, 9 February). Schireson, M. and Thakker, D. "The Money In Open-Source Software". Retrieved 19 December 2017 from <https://techcrunch.com/2016/02/09/the-money-in-open-source-software/>
- ²³ Lin, H. (2016, 2 October). "Attribution of Malicious Cyber Incidents: From Soup to Nuts". Columbia Journal of International Affairs. Abstract available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2835719
- ²⁴ CSO. Santarcangelo, M. (2016, 2 February). "Does attribution matter to security leaders?" Retrieved 12 December 2017 from <https://www.csoonline.com/article/3028907/leadership-management/does-attribution-matter-to-security-leaders.html>
- ²⁵ Microsoft. Charney, S. et al. (June 2016). From Articulation to Implementation: Enabling progress on cybersecurity norms. Retrieved 20 December 2017 from http://mscorpmedia.azureedge.net/mscorpmedia/2016/06/Microsoft-Cybersecurity-Norms_vFinal.pdf
- ²⁶ FireEye. (2017, 14 March). "FireEye Releases Mandiant M-Trends 2017 Report". Retrieved 20 December 2017 from www.fireeye.com/company/press-releases/2017/fireeye-releases-mandiant-m-trends-2017-report.html

- ²⁷ Lexology. Hogan, L. (2014, 18 April). "DOJ and FTC clarify antitrust implications of cybersecurity information sharing". Retrieved 21 December 2017 from <https://www.lexology.com/library/detail.aspx?g=2ab5ddc4-9791-44d7-b117-460ed5a1b3de>
- ²⁸ U.S. Department of Homeland Security. (2016, 21 June). "Automated Indicator Sharing (AIS)". Retrieved 12 December 2017 from <https://www.dhs.gov/ais>
- ²⁹ TechTarget. (n.d.). "botnet". Retrieved 21 December 2017 from <http://searchsecurity.techtarget.com/definition/botnet>
- ³⁰ Krebs on Security. (2012, 19 September). "Malware Dragnet Snags Millions of Infected PCs". Retrieved 12 December 2017 from <https://krebsonsecurity.com/tag/3322-org/>
- ³¹ Schneier on Security. (2016, 10 October). "Security Economics of the Internet of Things". Retrieved 12 December 2017 from https://www.schneier.com/blog/archives/2016/10/security_econom_1.html
- ³² Wired. Newman, L. H. (2016, 9 December). "The Botnet That Broke the Internet Isn't Going Away". Retrieved 12 December 2017 from <https://www.wired.com/2016/12/botnet-broke-internet-isnt-going-away/>
- ³³ TechTarget. (n.d.). "ISP (Internet service provider)". Retrieved 21 December 2017 from <http://searchwindevelopment.techtarget.com/definition/ISP>
- ³⁴ Upturn. (n.d.). Retrieved 12 December 2017 from <https://www.teamupturn.org/reports/2016/what-isps-can-see>
- ³⁵ Ars Technica. Anderson, N. (2007, 26 July). "Deep packet inspection meets 'Net neutrality, CALEA". Retrieved 12 December 2017 from <https://arstechnica.com/gadgets/2007/07/deep-packet-inspection-meets-net-neutrality/>
- ³⁶ Matania, E., Yoffe, L. and Mashkautsan, M. (2016). "A Three-Layer Framework for a Comprehensive National Cybersecurity Strategy". *Georgetown Journal of International Affairs*, 17(3), 77-84. doi:10.1353/gja.2016.0038
- ³⁷ World Economic Forum. (January 2011). *Personal Data: The Emergence of a New Asset Class*. Retrieved 21 December 2017 from http://www3.weforum.org/docs/WEF_ITTC_PersonalDataNewAsset_Report_2011.pdf
- ³⁸ Wired. Greenberg, A. (2014, 25 November). "Hacker Lexicon: What Is End-to-End Encryption?" Retrieved 12 December 2017 from <https://www.wired.com/2014/11/hacker-lexicon-end-to-end-encryption/>
- ³⁹ CSO. Korolov, M. (2017, 29 September). "Is universal end-to-end encrypted email possible (or even desirable)?" Retrieved 21 December 2017 from www.csoonline.com/article/3224410/encryption/is-universal-end-to-end-encrypted-email-possible-or-even-desirable.html
- ⁴⁰ National Institute of Standards and Technology (NIST), U.S. Department of Commerce. (n.d.). "Post-Quantum Cryptography, Project Overview". Retrieved 12 December 2017 from <https://csrc.nist.gov/Projects/Post-Quantum-Cryptography>
- ⁴¹ The Washington Post. Barnes, R. (2017, 16 October). "Supreme Court to consider major digital privacy case on Microsoft email storage." Retrieved 21 December 2017 from www.washingtonpost.com/politics/courts_law/supreme-court-to-consider-major-digital-privacy-case-on-microsoft-email-storage/2017/10/16/b1e74936-b278-11e7-be94-fabb0f1e9ffb_story.html?utm_term=.938c7d4f029c
- ⁴² Albright Stonebridge Group. (September 2015). *Data Localization: A Challenge to Global Commerce and the Free Flow of Information*. Retrieved 21 December 2017 from www.albrightstonebridge.com/files/ASG%20Data%20Localization%20Report%20-%20September%202015.pdf
- ⁴³ Computerworld. Thibodeau, P. (2011, 3 June). "Apple, Google, Facebook turn N.C. into data center hub". Retrieved 21 December 2017 from www.computerworld.com/article/2508851/data-center/apple--google--facebook-turn-n-c--into-data-center-hub.html
- ⁴⁴ Information Technology & Innovation Foundation (ITIF). Cory, N. (2017, 1 May) "Cross-Border Data Flows: Where Are the Barriers, and What Do They Cost?" Retrieved 21 December 2017 from <https://itif.org/publications/2017/05/01/cross-border-data-flows-where-are-barriers-and-what-do-they-cost>
- ⁴⁵ TechTarget. (n.d.) "data breach". Retrieved 21 December 2017 from <http://searchsecurity.techtarget.com/definition/data-breach>
- ⁴⁶ Center for Cyber & Homeland Security, The George Washington University. (October 2016). *Into the Gray Zone: The Private Sector and Active Defense against Cyber Threats*, p. 10. Retrieved 21 December 2017 from <https://cchs.gwu.edu/sites/cchs.gwu.edu/files/downloads/CCHS-ActiveDefenseReportFINAL.pdf>
- ⁴⁷ World Economic Forum confidential interviews, June-November 2017
- ⁴⁸ Carnegie Endowment for International Peace. Hoffman, W. and Levite, A. (2017, 14 June). *Private Sector Cyber Defense: Can Active Measures Help Stabilize Cyberspace?* Retrieved 21 December 2017 from <http://carnegieendowment.org/2017/06/14/private-sector-cyber-defense-can-active-measures-help-stabilize-cyberspace-pub-71236>
- ⁴⁹ Mercatus Center, George Mason University. Glosson, A. D. (August 2015). "Active Defense: An Overview of the Debate and a Way Forward". Retrieved 21 December 2017 from <https://www.mercatus.org/System/Files/Glosson-Active-Defense.pdf>
- ⁵⁰ Simmons & Simmons, elexica. (2014, 2 July). *Cyber risk and the extent of cover for "legally insurable" fines*. Retrieved 21 December 2017 from www.elexica.com/en/legal-topics/insurance/24-cyber-risk-and-insurability-of-fines
- ⁵¹ Bank of England. (2017, 5 July). "Cyber insurance underwriting risk". Retrieved 21 December 2017 from www.bankofengland.co.uk/prudential-regulation/publication/2017/cyber-insurance-underwriting-risk-ss
- ⁵² Akerlof, G. A. (August 1970). "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism". *The Quarterly Journal of Economics*, 84(3), pp. 488-500. JSTOR. Available from www.jstor.org/stable/1879431



COMMITTED TO
IMPROVING THE STATE
OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum
91–93 route de la Capite
CH-1223 Cologny/Geneva
Switzerland
Tel.: +41 (0) 22 869 1212
Fax: +41 (0) 22 786 2744
contact@weforum.org
www.weforum.org

World Economic Forum LLC
3 East 54th Street, 18th Floor,
New York, NY 10022, USA
Tel.: +1 212 703-2300
Fax: +1 212 703-2399
contact@weforum.org
www.weforum.org