HOW TO BAKE CYBERSECURITY REGULATIONS: INGREDIENTS FOR BETTER RESULTS

BY
MICHAEL DANIEL

Michael Daniel is the President & CEO of the Cyber Threat Alliance, a non-profit threat sharing organization for the cybersecurity industry. Prior to this position, Michael served for four and half years as President Barack Obama's cybersecurity advisor and U.S. Cybersecurity Coordinator.

Visit www.competitionpolicyinternational.com for access to these articles and more!

HOW TO BAKE CYBERSECURITY REGULATIONS: INGREDIENTS FOR BETTER RESULTS

By Michael Daniel

In most countries and economic sectors, organizations have traditionally faced few cybersecurity regulations. However, as the cybersecurity threat has worsened and the dependence on IT has grown, nations are increasingly turning to regulation as a method to improve their security. Yet, implementing effective regulations is not easy and governments could easily cause more harm than good. This article lays out five principles governments should follow to create more effective regulations: creating standards of care that vary by industry, criticality, and size; limiting complexity in any regulations; reallocating the security burden to the organizations in the ecosystem best positioned to handle it; avoiding zero-tolerance for failure; and harmonizing the rules across industries and jurisdictions whenever possible. Following these principles would produce regulations more likely to achieve the desired outcome of a more secure digital ecosystem.

01
HOW DID WE GET HERE?

For many years, the U.S. and most other Western nations left cybersecurity almost entirely up to market forces. In the late 1990s and early 2000s, that approach made sense, as networked information technology was just starting to make a meaningful difference in the economy. Adversaries were relatively few, and the downsides from security lapses were limited in scope for most organizations. To the extent policy makers talked about the Internet and cyberspace, the discussion typically focused on how the Internet could promote freedom and democratic values around the world.

Now the situation is radically different. We are connecting every kind of device imaginable to the Internet, from personal electronics to industrial control systems. Online platforms and services dominate have taken dominant positions in our economy. Virtually all significant economic and social activity now operates through cyberspace in some fashion. Governments use cyberspace to conduct activities ranging from diplomacy to service provision. The potential targets for malicious cyber activity have expanded enormously, creating a virtually unlimited supply of potential victims.

The threat has evolved too. Malicious cyber activity is no longer the province of sophisticated nation-state programs or socially disaffected people wearing hoodies. The criminal industry has diversified, specialized, and professionalized, becoming an enormously profitable business with low barriers to entry and very little risk. A few nations have found it in their interest to provide safe haven for these criminals, inhibiting other governments' ability to arrest and prosecute them. Nation-state adversaries are taking more risks, no longer shrink from exposure, and are willing to use cyber effects to cause disruption or destruction in the physical world. Hacktivists seek to get their message out through malign activities online.

Further, the ramifications of accidents or malicious activity have changed. We have become reliant on our networked technology at an individual, organizational, and society-wide level. Incidents that 20 years ago would have been minor annoyances can now generate catastrophic impacts: it is one thing if your spreadsheet crashes, but another thing entirely if your connected car crashes. Companies have suffered significant losses due to ransomware[2] and theft of intellectual property.[3] If a major cloud service provider suffers an outage, many of its customers have to suspend operations. The economic damages from a sustained disruption to cloud services would be enormous,[4] let alone the public health and safety impacts.

Finally, few people argue that the current level of cybersecurity is sufficient. Technology companies ship software with known flaws, organizations leave misconfigurations in place for years, and people clicking on links can bring down whole networks. The cybersecurity industry has grown to more than 3,000 companies offering a bewildering array of products, but users have difficulty determining whether using those products makes an organization more secure.[5] Most leaders see the risks as increasing, despite the enormous sums going into cybersecurity.[6]

> *"The threat has evolved too. Malicious cyber activity is no longer the province of sophisticated nation-state programs or socially disaffected people wearing hoodies*

In this context, where dependence is high, the risk large, bad activity rampant, and the defensive shortcomings painfully obvious, cybersecurity regulation is inevitable. No society can live indefinitely with the level of vulnerability we face today. Even countries with a strong anti-regulatory bias, like the United States, will impose additional regulations to reduce their vulnerability and increase security. Of course, some sectors, such as financial services, have had cybersecurity regulations for years, but the current environment will drive most democratic governments to expand cybersecurity regulations and to seek authority to regulate sectors where the government currently lacks such authority.

2 Combating Ransomware: A Comprehensive Framework for Action," Ransomware Task Force, Institute for Security and Technology, April 2021. https://securityandtechnology.org/ransomwaretaskforce/.

3 Emily Mossburg, J. Donald Fancher, and John Greline, "The hidden costs of an IP breach," Deloitte Review, Issue 19, 2016.

4 Robert Stevens, "Almost Everything is in the cloud and experts are worried," Fortune, October 24, 2022. https://fortune.com/2022/10/24/business-in-the-cloud-oxford-digital-economies/.

5 Research Report: "Cybersecurity Technology Efficacy: Is Cybersecurity the New Market for Lemons?" Debate Security, October 2020.

6 Global Cybersecurity Outlook 2023, World Economic Forum, Insight Report, January 2023. https://www.weforum.org/reports/global-cybersecurity-outlook-2023.

The question, therefore, is not whether cybersecurity will be regulated, but how it will be regulated. Asking the right question is important because it would be very easy for governments around the world to regulate cybersecurity stupidly. Cyberspace does not operate like the physical world and IT systems are not widgets. Many companies and organizations operate globally, cutting across multiple jurisdictions. If we try to regulate cybersecurity the way we have regulated other industries in the past, not only could we fail to improve cybersecurity, but we could make the situation worse. How can governments avoid falling into that trap?

The answer lies in approaching cybersecurity regulation with a clear set of principles focused on outcomes, rather than compliance, and taking into account how the Internet economy functions today. This article lays out five such principles for regulating cybersecurity smartly and avoiding the efficacy trap. Smart regulations would move us toward a safe, reliable digital ecosystem that supports national security, economic prosperity, and public health and safety while enabling innovation. A digital ecosystem with these characteristics would benefit everyone.

02
STANDARDS OF CARE LAY THE FOUNDATION

If governments want to regulate the security of the digital ecosystem effectively, they must establish the standards of care to which they will hold organizations accountable.[7] Businesses and other organizations need clear rules to determine whether they need to make changes in their policies, where they should invest resources, and how they should prioritize activities. Customers need standards to make purchasing decisions, and courts need them to make decisions about liability or negligence. Without standards of care, regulations can't be effective. Industry has not settled on such standards of care voluntarily, and the lack of agreed upon "standards of care" has greatly inhibited the ecosystem from improving its security over the past 15 years.

Two reasons are typically cited for why such standards have not emerged. One argument is that cybersecurity experts cannot agree on what standards to enforce.[8] Another is that IT changes too rapidly for standards of care to keep up.[9] While these arguments held some weight in the past, they are no longer accurate. The cybersecurity community knows what policies, practices, and structures increase security and which ones do not, even if it has trouble measuring exactly how much difference any given policy or practice makes. Examining the various control lists cited in the footnote reveals more similarities than differences. Moreover, these "best practices" have been consistent for at least a decade. For example, security experts have known since the early-2000s that a username and password alone does not provide sufficient security; a second method or "factor," such as a text message or authentication application, is needed. We will not return to a world where only passwords are sufficient. Another example is the National Institute of Standards and Technology's Cybersecurity Framework.[10] While some of the informative references have evolved since its publication in 2014, the core Framework and its elements have remained stable for over nine years. Therefore, adopting standards of care that will not become rapidly obsolete is not only feasible but long overdue.

Of course, the full standards of care that a company should follow will vary by industry, size, and function. A single standard of care will not work for the entire economy. Some entities are more systemically important than others, technologies differ across industries, and available resources vary wildly between small businesses and large enterprises. These factors should lead governments to avoid trying to set one overarching "cybersecurity standard of care" for an entire country; instead, tailored standards will be needed for different sectors and contexts.

However, if the *what* won't change rapidly, the *how* will. While SMS texts might still be a reasonable second factor in 2023, many malicious actors are rapidly learning how to spoof or bypass that technology. In fact, state of the art authentication approaches eliminate passwords altogether. As a result, effective regulations will avoid specifying the means that organizations use to achieve a security outcome. Focusing on results without mandating specific methods will produce far better outcomes and retain sufficient flexibility to adopt new technologies as they emerge.

7 In this context, a standard of care is the level of cybersecurity a reasonable person would expect an organization to employ, based on the industry, company size, and other factors.

8 For example, the Center for Internet Security has its Critical Security Controls (https://www.cisecurity.org/controls), BitSight has a list (https://www.bitsight.com/blog/cybersecurity-controls-types), and the UK government has its cyber essentials program (https://www.ncsc.gov.uk/cyberessentials/overview).

9 "The Law Can't Keep Up with New Tech. Here's How to Close the Gap," Daniel Balan, World Economic Forum Blog, June 21, 2018. https://www.weforum.org/agenda/2018/06/law-too-slow-for-new-tech-how-keep-up/.

10 https://www.nist.gov/cyberframework/.

Finally, much of the expertise needed to set standards of care resides in the private sector. Calibrating standards of care will involve both cybersecurity and domain knowledge, along with legal and process knowledge. Few government regulators will have all the necessary knowledge in-house. Therefore, an open, inclusive process will generate far superior outcomes over approaches that rely solely on government staff.

03
COMPLEXITY IS THE ENEMY

In cybersecurity, complexity is almost always an enemy. Managing IT hardware and software, on-line services, and data in a small or medium enterprise is challenging, let alone in a large multinational enterprise. Complex environments inevitably create gaps and overlaps, which malicious actors exploit to carry out their activities. Therefore, a long-established best practice in cybersecurity is to reduce complexity.[11]

Similarly, governments should follow this principle in developing regulations. Keeping cyber regulations simple and focused on a limited number of outcomes will improve security far more than long compliance checklists. Since the standards of care will have to vary among sectors to some degree, keeping the regulations simpler will help companies that operate in multiple sectors manage that complexity. Finally, many companies operate in multiple jurisdictions. Creating a plethora of different compliance requirements will simply divert resources from security functions and reduce efficacy.

Another reason to keep regulations simple is that we still have difficulty measuring cybersecurity. While we know the practices that produce more secure outcomes, evaluating those improvements objectively or precisely remains difficult. How much more secure does proper network segmentation make a company? What amount of time to detect an intruder is the acceptable minimum? How many vulnerabilities does a company need to patch each month? The cybersecurity industry cannot answer these questions very well yet. Given these limitations, it is foolish to pretend we have more knowledge than we do. Therefore, governments should embrace simplicity in developing cybersecurity regulations and assume that less will generate better outcomes, at least for the foreseeable future.

04
ALLOCATING THE SECURITY BURDEN EFFICIENTLY IS NECESSARY

Currently, information and operational technology markets follow an unusual structure. Software producers bear little to no liability for product performance or for any exploitable vulnerabilities in them. Service providers have few security obligations from a legal standpoint in most jurisdictions, nor do platform providers have to guarantee any level of security. Instead, the security burden lies almost entirely on the end-user, whether that's an individual, small or medium business, large enterprise, or a huge multinational corporation. Few products and services have this liability structure.

This unusual market structure emerged as information technology and the Internet developed in the 1980s and 1990s. Originally, this structure helped fuel innovation, allowed for rapid deployment and adoption, and created minimal downsides. When software crashed or networks went down, most organizations suffered limited damage or disruption and could recover quickly. Individuals mostly lost unsaved data.[12] Software makers could issue patches for vulnerabilities when they emerged, and the volume of patches was small enough that users could theoretically keep up, even if most didn't in practice. Adversaries could not cause systemic harm in this environment. The Internet fostered connectivity in previously unobtainable ways, generating massive changes that were mostly perceived as beneficial. As a result, few policy makers or users objected to these market dynamics.

However, in the early 21st century, this market structure has resulted in several systemic problems that hinder effective cybersecurity. First, a lack of penalties for security flaws reduces the incentive to spend time and money on building security into software; coupled with economic incentives that strongly favor being "first to market" over "secure to market," even software vendors who want to prioritize security face an uphill battle. As a result, the number of software vulnerabilities has become enormous and elaborate systems have evolved to help identify, publish, rate, and catalog known software vulnerabilities. For example, MITRE published 25,068 unique vulnerabilities during the 2022 calendar year alone, a 24.3 percent increase over 2021 and an average of 69 vulnerabilities *per day*.[13] Based on research by the Cyentia Institute and Kenna Security, the

11 The Digital Big Bang, Phil Quade, John Wiley & Sons, 2019.

12 Although many of us remember the frustration of losing an almost finished college paper.

13 https://www.cve.org/.

average monthly remediation rate for open vulnerabilities in 2022 was about 15.5 percent, meaning that most organizations are falling ever further behind in patching their security holes.[14] Even the most well-resourced organizations cannot keep up with this volume.

Second, the main demand signal from users is convenience and anything that detracts from a convenient experience potentially reduces business. Most security involves some cost in terms of time or mental capacity, which often places security and convenience at odds. Therefore, vendors and on-line service providers have little incentive to consider security in their default settings and configurations. The result is that security features usually require user action to enable them, and configurations are set to their most open state; since people often do not adjust these settings, software and services are often more vulnerable than they need to be.

Third, service providers are incentivized to deliver all information to end-users as quickly as possible, even if it contains identifiable, well-known malware or denial of service packets. Conversely, if a service provider tried to filter out known bad traffic, they lack protection for the inevitable good-faith false positives that might impact a customer. End-users can often pay for filtering services from managed security service providers or content providers, but this structure again requires users to take action to protect themselves.

Thus, the current market structure places almost the entire security burden on the part of the ecosystem least capable of handling it from a technical, organizational, and efficiency point of view – end-users. Trying to improve security through end-users is the least efficient approach with the lowest leverage, so it is not surprising that our efforts have produced mixed results. Instead, if we reallocate the security burden to the more capable elements in the ecosystem, then both efficiency and coverage will increase. What would such a shift look like?

One possible change would be to require software producers and platform providers to use more secure coding practices to reduce the number of vulnerabilities created as software is developed. Such a mandate would provide the "cover" private sector companies need to justify slowing their development processes to incorporate better coding practices and security reviews. Many companies want to implement these changes, but market pressures make such a choice prohibitive. Governments are also considering whether to use their procurement authority to require government vendors to reduce the number of vulnerabilities affecting security in products acquired by the government.[15] Since developers would not use one process for the government and another for other customers, this requirement would reduce the known exploitable vulnerabilities in software.

> ***Currently, information and operational technology markets follow an unusual structure. Software producers bear little to no liability for product performance or for any exploitable vulnerabilities in them***

As a second adjustment, governments could require most IT products to utilize secure configurations and settings as the default. This approach would rely on well-established principles of behavioral economics to achieve a better outcome. For example, we know that if employees are automatically enrolled in retirement systems, fewer than 5 percent will opt out; however, if employees have to make the effort to enroll, only about 50 percent will do so.[16] Similarly, security should be opt out, not opt in. We know that most people do not adjust the default settings; if the product has the more secure settings "out-of-the-box," then we would face far fewer configuration issues.

Other ways to reallocate the security burden exist. For example, governments could renegotiate the social "contract" with Internet Services Providers to have them filter out known malicious content by default. To be effective, this approach would require answering important questions such as how to designate malicious content and who would identify it; establishing such mechanisms would be highly challenging for democratic societies, but it might be possible to construct methods where the benefits outweigh the costs and risks. While people can and should argue about the benefits and costs of specific approaches for reallocating the security burden, the underlying point is that the current allocation does not reflect an immutable law of nature, but rather policy choices that can be changed to benefit security.
This principle comes with three caveats. Some burden must

14 Cyentia Institute and Kenna Security. 2022. Prioritization to Prediction Vol 8. (2022). https://www.kennasecurity.com/resources/prioritization-to-prediction-reports/.

15 See section 6722 of the House-passed version of the National Defense Authorization Act of 2023. https://www.congress.gov/bill/117th-congress/house-bill/7900/text.

16 John Beshears, James Choi, David Laibson & Brigitte C. Madrian, "The impact of employer matching on savings plan participation under automatic enrollment," in David A. Wise, ed., Research findings in the economics of aging (Chicago, IL: University of Chicago Press, 2010), pp. 311–327

remain on end-users, both for practical reasons and to reduce moral hazard. Just as auto manufacturers cannot be held liable for someone driving recklessly, software producers cannot and should not be held liable for users behaving recklessly in cyberspace, such as turning off firewall protections. Even with all the safety features built into cars, we still expect drivers to put on their seatbelts and not text while driving. Second, we cannot expect these changes to be "free;" for example, if societies want to shift some of the security burden to telecommunications providers, then those companies must receive increased legal protection and compensation for their activities on behalf of society. Users will have to live with some cost to convenience. Finally, we should not expect fundamental market restructuring to happen quickly. This shift can only occur over several years, perhaps a decade.

05

ZERO-TOLERANCE WON'T WORK

Many software producers and platform providers rightly worry about the flip side of the burden allocation question. Regulation frequently takes a binary, compliance form – either you meet this checklist, or you do not, and the regulator has zero-tolerance for non-compliance. Such a zero-failure mindset will not work in cybersecurity.

First, we don't know how to write perfectly secure code and writing significantly more secure code usually involves considerable effort.[17] In fact, perfectly secure code is likely physically impossible to achieve. Since all software will therefore have some number of exploitable vulnerabilities in it, holding companies to a "no flaws ever" standard will fail. Instead, the focus should be on incentivizing the use of techniques proven to minimize the creation of exploitable vulnerabilities and having a robust remediation process for when new ones are discovered. It should also involve methods for switching out unsupported code, adopting modern operating systems, and updating open-source libraries.

Second, market economics favor using general-purpose computing chips versus special-purpose chips. Even ostensibly special-purpose devices usually use general-purpose chips because such chips are generally the cheapest. By definition, a general-purpose chip can be made to do a wide variety of tasks and functions, including those never envisaged by the original designer or current user; that's why, for example, an internet-connected camera can be incorporated into a botnet. Research has shown that restricting the functionality of a general-purpose chip is effectively impossible.[18] Therefore, malicious actors will always be able to find ways to make almost all current IT devices perform an undesired task or function, no matter how secure the code running on it. Absent a change in market dynamics toward favoring special-purpose chips, this factor will limit our ability to eliminate malicious activity.

Third, we face intelligent adversaries who actively seek to circumvent security. They adapt, evolve, and improve over time, learning from each other and defenders' actions. The incentives for malicious actors are very strong – enormous sums are at play for criminals and substantial national security advantages in the nation-state arena. No matter how effective defense becomes, at least some number of malicious actors will continue their activities. We cannot expect those adversaries to fail to achieve their goals every time.

Fourth, human users are fallible. Adversaries already exploit human psychology through phishing and similar techniques to gain access to networks. We will not eliminate human mistakes, such as clicking on a malicious link or succumbing to a scam or even being bribed to become an insider threat. Some percentage of these attempts will succeed, no matter how effective the technology gets.

Effective cybersecurity regulations will acknowledge these factors and not try to achieve a zero-failure outcome. Instead, in addition to reasonable standards of care, they should require companies to have the ability to respond to breaches or business disruption. If a company meets the standards of care and it executes an acceptable incident response plan when an incident occurs, then it should not be punished or reprimanded – instead, it should be held up as an example of what effective cybersecurity looks like.

17 Is it even possible to be "completely secure"? | by April Wright | Medium, https://medium.com/@aprilcwright/is-it-even-possible-to-be-completely-secure-6c7a92a297a9.

18 Presentation by Thomas Dullien at CyCon X, Tallinn, Estonia, May 2019.

06

HARMONIZATION DETERMINES HOW SMOOTHLY REGULATIONS OPERATE

Another legitimate concern about increased cybersecurity regulation revolves around fragmentation. For any organization conducting business in more than one jurisdiction, fragmentation in cybersecurity regulations could mean not only having to comply with multiple regimes, but potentially conflicting ones as well. Such compliance fragmentation would make it difficult for companies to make efficient, consistent investments in cybersecurity and it would increase the likelihood of opening new gaps and seams that adversaries will exploit. It would also create an overhead burden that reduces the resources available for cybersecurity activities. Finally, complying with hundreds of different regulatory regimes will not be physically or economically feasible, which either means that a company simply takes the risk of ignoring some of them or they exit a market. Neither result is desirable. As a result, regulatory fragmentation threatens to undermine the goal of increased security.

Of course, not all regulatory differences are bad. Differing regimes can reflect varying risk tolerances or policy choices, which are a nation's prerogative. Dealing with those variations is the cost of doing business in multiple countries. However, different regulatory regimes more typically occur due to bureaucratic tendencies than genuine policy differences. Sovereign nations understandably like to set their own policies, and regulators within a given country like to develop regulations of their choosing. Yet, these bureaucratic tendencies often result in differences of style or format without difference in substance, or even worse, conflicting requirements that are physically impossible to comply with at the same time.

Therefore, successful regulatory regimes will minimize fragmentation. They will recognize that most software manufacturers, service providers, and platform providers operate in many different jurisdictions and all platform providers essentially operate globally. Effective regimes will also acknowledge that critical infrastructure companies in many industries often span multiple markets. They will seek to harmonize their requirements wherever possible and ensure that differences reflect distinct policy choices. We have models for such harmonization in industries like air transportation, so this concept is not new in the regulatory space. We need to apply it to cybersecurity.

07

CONCLUSION

Cybersecurity regulation may be inevitable, but smart, effective regulation is not. Such an outcome is achievable, however. Despite popular belief, government agencies learn, adapt, and change their practices over time. Therefore, if cybersecurity experts engage with governments as they develop cybersecurity regulations, the resulting rules have a much greater likelihood of producing the desired outcome while minimizing the unintended consequences. Using the principles laid out above would create a digital environment where cybersecurity is far easier for most people and organizations, expectations are clearer, and malicious actors are much less successful. That's a goal worth investing in.

Of course, not all regulatory differences are bad. Differing regimes can reflect varying risk tolerances or policy choices, which are a nation's prerogative

CPI SUBSCRIPTIONS

CPI reaches more than **35,000 readers** in over **150 countries** every day. Our online library houses over **23,000 papers**, articles and interviews.

Visit **competitionpolicyinternational.com** today to see our available plans and join CPI's global community of antitrust experts.

CPI COMPETITION POLICY® INTERNATIONAL