30 October 2023

SIPA Cyber
Columbia University SIPA
420 W118th Street, Suite 1337
New York, NY 10023

Office of the National Cyber Director
Executive Office of the President
1600 Pennsylvania Avenue, NW
Washington, DC 20500

Comments on Cyber Regulatory Harmonization (ONCD_FRDOC_0001-0002)

We commend the Office of the National Cyber Director seeking input on new regulations and regulatory harmonization. The National Cybersecurity Strategy (NCS) correctly assesses that cyber regulation is needed as markets have failed to deliver security. However, the government and outside researchers still have substantial homework to do to ensure regulations are correctly targeted for the best chance of success and harmonization, with the least cost of compliance.

Accordingly, we at the School of International Affairs at Columbia University in in the process of standing up a new Cyber Regulations Lab, part of our ProjectNCS, partially funded by an unrestricted gift from Google. ProjectNCS will also examine the core assessment question of "are we winning?" That is, we will assess indicators if cyberspace is becoming more defensible, as called for in the strategy. This effort complements our New York Cyber Task Force – which helped develop the concept of a more defensible cyberspace, which is foundational to the NCS – and our research and convening on cyber risk to financial stability, in partnership with the Federal Reserve Bank of New York.

Our regulation effort is still new and so cannot engage with the range of detailed questions in the ONCD Request for Information. It will focus instead on three areas: the need for a new regulatory strategy to drive harmonization and other priorities, the need to examine market failures and full range of public-policy tools, and a framework to better understand when to use "performance-based" or other flavors of regulations.

That final section goes into substantial depth as the cyber-regulatory community's use of this performance-based regulation is nearly the opposite of how it is used in regulation of physical assets. Performance-based regulations are generally more difficult to harmonize than those mandating general principles or requiring specific management processes.

Recommendation 1: ONCD Should Lead Development of a New Cyber-Regulation Strategy

Regulation is the most important, complex, and politically sensitive project ever undertaken by the Federal government. Unlike other topics, say workforce and education, ONCD has just one shot to get this correct.

ONCD should begin coordinating a new strategy to plan, integrate, and execute cyber regulations across the government more comprehensively. This strategy, or perhaps a less formal roadmap, would slot underneath the NCS to coordinate the full range of regulatory work, from harmonization of existing

regulations, minimum security baselines for critical infrastructure, pushing for software liability, exploring options for regulations for platforms or major service providers, and continuing the needed homework mentioned in the following sections.

These are critical and complex issues which touch the heart of the role of government in a free society. They deserve a single strategy, signed off by deputies or principals.

A regulatory strategy or roadmap could be substantially harder to coordinate than the recent strategy on cyber education and workforce, as so many of the regulatory agencies involved are independent and cannot be bound by such a document. Moreover, the year before a new president election may not be the best time for such a politically sensitive and difficult strategy.

Fortunately, the Forum for Independent and Executive Branch Regulators, led and newly reinvigorated by Jessica Rosenworcel, the chair of the Federal Communications Commission, has already been making strong progress. The right process for a strategy or roadmap should reinforce rather than compete with their efforts.

Recommendation 2: ONCD, SRMAs, and Other Stakeholders Should Determine Market Failures and Corresponding Regulatory Toolkits

Regulations cannot be harmonized or "tailored for each sector's risk profile," unless the government better understands which sectors of the economy suffer from which specific market failures.

The original White House cyber policy document, PDD-63 of 1998, specified that "regulation will be used only in the face of a material failure of the market to protect the health, safety or well-being of the American people." The NCS repeatedly asserts that markets have indeed failed and so regulation is now required.

Yet there is surprisingly little policy or academic work on which markets are failing and how. Other than for software liability, which is well studied, too much of the remaining work relies on analogy, anecdote, and gut instincts.

ONCD should work with the National Economic Council and the Cybersecurity and Infrastructure Security Agency and other Sector Risk Management Agencies to develop official assessments. The Offices of the Chief Economist at SRMAs should have a leading role. Meanwhile academic efforts, like those with the Cyber Regulations Lab here at Columbia University, will continue our research on these issues.

For example, the commercial nuclear sector likely has high negative externalities (a lot of people suffer if they fail) and market power (there are only a few operators) but low information asymmetry and public goods (the regulators have excellent knowledge of their risks and operations and can raise rates if necessary).

By comparison, the water and wastewater sector may have severe market failures across the board: many people will suffer, there is usually one provider, few externally know their internal risks, and it is hard to exclude those that won't pay for more security. While minimum cybersecurity baselines might be similar between these two sectors, these sectors will otherwise need tailored regulations.

A related workstream should create a toolkit of public-policy remedies for each kind of market failure. When financial regulators assess a bank is too big to fail (a systemically important financial institution, in the lingo), they have a defined set of tools to address that market failure before it becomes acute: requirements for additional capital, additional stress tests, and living wills to simplify their breakup, should that become necessary.

The toolkit for cyber policymakers will not be so straightforward but there are some obvious candidates. If a sector suffers from information asymmetries, for instance, it is appropriate to regulate for increased transparency such as with mandatory cyber incident reporting or software bills of materials. Minimum cybersecurity baselines should be used to reduce negative externalities.

There is a limited list is such public-policy tools. Once each is associated with the associated market failure, and there is a mapping of which sectors have which market failures, it will be far simpler to determine which set of tools are the best fit – and why – for each sector. This will substantially ease efforts to harmonize regulations.

Recommendation 3: To Harmonize, ONCD and Other Stakeholders Should Better Align from Existing Regulatory Frameworks

Companies in critical infrastructure sectors must patch their systems to prevent cyber incidents from disrupting their operations and avoid cascading impacts on the U.S. economy and national security. But which one of these regulations is better at achieving that goal?

1. "Identify, report, and correct information and information system flaws in a timely manner. "
2. "For patches and updates that are listed on CISA's Known Exploited Vulnerabilities Catalog and have a NIST Base Score of "Critical" … the patch/update must be installed within 15 days of its availability. All other updates and patches must be installed within 30 days of availability."

The first, from the U.S. Department of Defense, provides regulated entities maximum flexibility to implement the broad principles. But such vague guidance offers no suggestions on how serious a flaw would have to be to be fixed and just how long is timely. A company might patch the flaw in a month only to be punished because it wasn't done in a week.

The second, from a now superseded U.S. regulation for pipeline operators, is strikingly specific. Entities won't have much doubt about the regulator's expectations, which could benefit less-mature cyber teams or those who don't want to be punished after the fact by a crusading regulator. But many other companies would feel handcuffed. Because it doesn't treat all vulnerabilities as the same, this regulation is risk-based, but that risk assessment is being made by others and not the entities themselves. Just because a vulnerability is rated as "critical" by someone else, doesn't mean it is critical to them. A company which successfully met these rigorous deadlines might have not have any resources to accomplish anything else.

The NCS puts such issues at the center of the cybersecurity debate by pushing for new regulations which should be "performance-based" and agile as well as tailored for each critical-infrastructure sector and "harmonized to reduce duplication."

But what makes a regulation performance-based? Are they always the best choice or does tailoring for each sector mean some regulations should be something other than performance-based? Are performance-based regulations easier to harmonize across regulators and jurisdictions?

There is very little research to guide policymakers. Fortunately, Jim Dempsey set an initial path for such research in an influential Lawfare article.

This third section (derived from a forthcoming article) takes further important steps down Dempsey's path, referencing past regulatory work to understand performance-based regulations; how they differ from rules-based, principles-based and other kinds of regulation; which existing cyber regulations fall in each such category; when regulators should use or avoid principles-based regulations; and under which circumstances, more generally, is one kind of regulation the clear, best choice.

As it turns out, of the two regulations on patching mentioned above, the second one – with specific and measurable outcomes– is by far the more performance-based. But such tightly prescriptive rules are likely not what the White House wanted to encourage. It was soon superseded after cries for more performance-based regulation when it was in fact one of the most performance-based cyber regulations ever. What critics wanted was more management- or principles-based regulations.

Framework of Cyber Regulations

The most important distinctions in the regulation of physical assets are the types of commands (specifying either ends or means) and the breadth of their focus (as micro or macro). These are all broadly command and control regulations, that is the government specifies what regulated entities should seek or avoid rather than relying on tax incentives or self-regulation.

Cary Coglianese – in the canonical article describing this framework – distinguishes that when regulators specify ***ends***, they direct targets to achieve or avoid certain outcomes: for example, 10 grams of ground allspice cannot have more than 30 insect fragments or 1 rodent hair. Or in cybersecurity terms, select systemically important financial firms had to "achieve recovery and resumption within two hours" after a disruptive event, according to a 2003 requirement.

Means, by comparison, regulate for some proxy for regulator's actual goals by mandating or forbidding particular behaviors, processes, or technologies. A common cybersecurity means regulation is some version of having a "24 × 7 × 365 computer incident response capability for cyber incidents," as is required for the chemical sector.

Micro-level regulations "require either the adoption of specific means or the attainment of concrete outcomes" while those at the **macro-level** require only the most general means or ends. For example, the Transportation Security Administration's regulation for the pipeline sector to classify their information-technology assets ("Review and assess pipeline cyber asset classification as critical or non-critical at least every 12 months") is more micro than their similar rule for rail ("Facility has an asset inventory of all critical IT systems"). Perhaps this difference is due to differences in the risk profile of each sector, perhaps not.

Table 1 below has an initial 2x2 framework to distinguish means-ends and micro-marco with representative examples from cyber regulations of the past two decades.

Table 1: Framework of Cyber Regulatory Controls

	Means Management-Based. Mandates the use of security technologies or behaviors	**Ends** Performance-Based. Achieve or avoid certain security outcomes
Micro Rules-Based, Specific	• CISO shall report in writing at least annually to the Covered Entity's board of directors (23 NYCRR Part 500) • Board or an appropriate board committee has cybersecurity expertise or engages experts (FFIEC Security Assessment Tool) • Designate and use a primary and at least one alternate Cybersecurity Coordinator (TSA Security Directive Pipeline-2021-01B) • Report [incidents] as soon as practicable, but no later than 24 hours after a cybersecurity incident is identified (TSA Enhancing Rail Cybersecurity) • Implement MFA for access to assets using the strongest available method for that asset (CISA CPG v.1.0.1) • Facility has a defined 24×7×365 incident response capability (CFATS RMBP-8) • Review and assess pipeline cyber asset classification as critical or non-critical at least every 12 months (TSA Pipeline Security Guidance) • Automated tools enable tracking, updating, asset prioritizing, and custom reporting of the asset inventory (FFIEC Security Assessment Tool) • Patch all KEV in all public facing systems (CISA CPG v 2022) • Provide a purchaser an SBOM for each product directly or by publishing it on a public website (EO 14028) • Include written procedures, guidelines and standards to ensure the use of secure development practices for in-house developed applications (23 NYCRR Part 500)	• Develop the capacity to recover and resume … activities within the business day on which the disruption occurs with the overall goal of achieving recovery and resumption within two hours after an event (Sound Practices Memo 2003)
Macro Principles-Based, General	• Establish technical or procedural controls for cyber intrusion monitoring and detection (TSA Pipeline Security Guidance) • Tools and processes are in place to ensure timely detection, alert, and activation of the incident response program (CRI Profile v1.2) • Have a cyber risk management framework that is reviewed and approved by the Board (CRI Profile v1.2) • The facility practices the concept of least privilege (CFATS RMBP-8) • Establish a vulnerability disclosure program (EO 14028) • The facility practices the concept of least privilege (CFATS RMBP-8) • The facility has an asset inventory of all critical IT systems (CFATS RMBP-8)	• The firm manages information and data consistent with its risk appetite and tolerance for disruption to protect the confidentiality, integrity, and availability of data and systems (Sound Practices 2020) • The organization mitigates cybersecurity incidents in a timely manner (CRI Profile v1.2) • Organizations are capable of safely and effectively recovering from a cybersecurity incident (CISA CPG v.1.0.1)

Categories of Regulation and When to Use Which

This simple 2x2 matrix simplifies the definition of performance-based regulation as compared to other types of regulation, allowing regulators to far more easily tailor regulations to the risks of their particular sector.

Performance-based regulations are micro-ends.

Also called outcome-based, these are requirements to *mandate or avoid or achieve the specific outcomes which are the ultimate concern to the regulator*.

Recovery-time objectives, like the two-hour limit mentioned above for systemically important clearing and settling firms, are performance based: regulators want a specific result (recovery and resumption of clearing and settling) in a specific timeframe (two hours).

Because they specify a regulator's desired outcomes, it is no wonder that the National Cybersecurity Strategy encourages this model. But a key characteristic of performance-based regulations is that they succeed "only if government agencies are able to specify, measure, and enforce performance," in Dempsey's words.

Unfortunately, cybersecurity is "immensely difficult to measure," because cybersecurity is not like allspice. Technology moves quickly, there are creative and motivated adversaries, and what worked yesterday won't ensure success tomorrow.

Recovery-time objectives are one of the few areas in cybersecurity where outcomes are so easily measurable. It is accordingly one of the few existing performance-based cyber regulations since, for all their strengths, such regulations are often an unsuitable – or at least unavailable – a choice for cybersecurity.

Even where measurement is possible, a performance-based regulation may decrease, not increase flexibility. As mentioned above, giving pipeline operators just fifteen days to patch critical vulnerabilities is measurable but far too specific. Patching at scale is complex and only a small fraction of vulnerabilities are ever exploited: such strict deadlines can undermine security.

Performance-based regulation is, according to Coglianese, most appropriate for sectors that are broadly homogeneous ("most firms have similar operations that remain stable over time"). They may be more suitable, for example, for cyber regulations for the pipeline sector than fintech.

To summarize: performance-based regulations are most suitable where outputs can be measured for mostly homogeneous, stable sectors and entities.

Management-based are the opposite of performance-based: macro-means.

These regulations mandate general planning and management practices which hopefully achieve the regulator's ultimate objectives, especially when outcomes can't be measured. As Columbia cybersecurity professor Steve Bellovin commented when reviewing a draft of this section, in cybersecurity "many things can't be measured, but some processes really, really help."

There are many such cyber regulations: "Establish technical or procedural controls for cyber intrusion monitoring and detection," as called for in pipeline regulations or "tools and processes are in place to ensure timely detection, alert, and activation of the incident response program," per the Cyber Risk Profile for the finance sector.

While such regulations can still be quite specific and mandatory, they are still more general than performance-based regulations. They are neither dictating specific technologies or controls, for example, nor specifying how serious an incident needs to be before it triggers a response.

"Management-based regulation is worth considering any time the government confronts hard-to-assess risks generated by many diverse firms," advises Coglianese. Accordingly, it may be a better fit for much of cybersecurity where technology is constantly changing.

Principles-based regulations are *macro-ends*.

A primary goal of principles-based regulations is to "give regulated entities flexibility in how best to achieve a desired objective," according to Heath Tarbert, then-head of the U.S. Commodity Futures Trading Commission, in 2020. They are macro-ends.

Compare these two regulations: "Tools and processes are in place to ensure timely detection, alert, and activation of the incident response program," and "the organization mitigates cybersecurity incidents in a timely manner." Both come from the same document, the Cyber Risk Profile for the finance sector and are macro. But the first regulates ends (successful, quick mitigation) and is *principles*-based while the second regulates means (tools and processes) and is accordingly *management*-based.

Tarbert noted several strengths of principles-based regulation:

- "Generally promotes a more flexible regulatory approach;"
- "Enhances the responsiveness of regulation to market innovation and other developments [and] can help 'future proof' regulatory requirements;"
- "Discourage 'loophole' behavior and 'checklist' style approaches to compliance with the law;"
- "Encourages the more direct involvement of C-suite or other senior management rather than relying on lower ranking compliance personnel;"
- "Also can lower compliance costs."
- Can encourage innovation by facilitating "the development of new business models, products, and internal processes;" and
- "Helps to promote comparability and convergence among international regulators. While different national regulators rarely agree on specific, granular rules … they nonetheless frequently can reach consensus on principles;"

This last strength makes principles- and management-based regulations particularly relevant for harmonization, given the White House's push for regulatory harmonization. It is far easier to harmonize global financial regulations when they are as broad as this rule by the Office of the Superintendent of Financial Institutions Canada: "timely response, containment and recovery capabilities are established."

Principles-based cyber regulations are less measurable than performance-based, since they are more general and each regulated entity is responsible for meeting the goals as they think best.

These regulations are a particularly good fit for sectors like finance which have mature teams able to translate the broad regulatory outcome into entity-specific behaviors and measure compliance. The result may still lead to compliance checklists, but at least they are checklists of the firm's own design.

Critically, principles-based regulation requires a robust system of external oversight, coupled with post-hoc enforcement of the rules, such as the extensive structure built for the finance sector. Proper cybersecurity oversight would require a substantial investment to create or expand government cyber-regulatory agencies or the creation of self-regulatory organizations, an idea proposed in class by Columbia students.

Rule-based regulations are the complement of *principles-based*: micro-level and either means or ends.

Rules-based regulation has a bad reputation as a mere compliance exercise that doesn't reduce risk, just ineffectual and endless checklists. But performance-based regulation, as called for in the National Cybersecurity Strategy, is a form of rule-based regulation.

An advantage of prescriptive rules, compared to principles, is that they can provide greater clarity with regulators or against private litigation. For example, clearing and settling firms may have thought a government-mandated recovery-time objective of two hours was unwarrantedly quick. But it is far clearer objective than that stated in the Cyber Risk Profile ("The organization mitigates cybersecurity incidents in a timely manner") or the requirement of the Canadian financial regulator (for "timely response, containment and recovery"). A firm may consider their eight-hour mitigation to have been exceptionally prompt, only to be retrospectively fined by a regulator with different expectations.

A rules-based approach, according to Tarbert, "is most effective when regulators are able to adopt rules that can endure." While it may sometimes seem that cybersecurity threats and technology are changing too quicky to set prescriptive rules, many existing cybersecurity regulations seem appropriately rules based:

- Applications running executable code should be disabled by default on all IT and OT assets to reduce the risk of malware (Cybersecurity Performance Goals, v 2022);
- The institution belongs or subscribes to a threat and vulnerability information sharing source(s) that provides information on threat (FFIEC Security Assessment Tool);
- The CISO of each Covered Entity shall report in writing at least annually to the Covered Entity's board of directors or equivalent governing body (NY Department of Financial Services, 23 NYCRR Part 500);
- The facility has an asset inventory of all critical IT systems (CFATS for the chemical sector);
- Use administratively separate build environments (EO 14028 for entities selling to the U.S. government).

Checklists may get demeaned but are an effective tool to ensure success on repetitive and complex tasks.

Moving Forward

The framework in Table 1 is only an early effort to apply the lessons from regulating physical assets, building on the earlier excellent work of Dempsey.

Even this early work is enough for several suggestions.

First, as they work to harmonize existing regulations and develop new ones, regulators should assess whether such rules are micro or macro, ends or means. As noted above, sometimes regulations in the same document are found in different quadrants of the matrix.

There can be good reasons for a single document to have regulatory controls in different quadrants: some might be more measurable (and so performance-based) while others are not and need to specify processes (management-based). But some controls are in different quadrants simply because until now there hasn't been a simple table with such quadrants. Either way, regulators should review their existing and proposed rules to determine which fit where and why.

Second, ONCD should work with SRMAs, regulators, and chief economists should use this framework to determine the best type of regulation for their sector, leveraging the assessment of market failures (Recommendation 2 above).

A performance-based (micro-ends) approach is generally the best fit only for rules where the output can be easily measured and for sectors which are relatively stable and one regulated entity looks much like the rest. This might be a fit for the water and waste-water sector. Regulations for the finance sector – fast-moving, heterogeneous, deeply experienced in tracking their own compliance, and regulated by many entities – should likely be based on principles, not rules.

Third, to maximize chances for harmonization, regulators should aim for principles-based regulations which are more amenable than those with more-specific rules. After principles are agreed upon, aligning the more-exact rules, if necessary, can come later with more experience and trust.

Again, we appreciate the opportunity to comment. In addition, thanks to Jim Dempsey, Josephine Wolff, and the participants at the 2023 Cybersecurity Law and Policy Scholars Conference at The Fletcher School at Tufts University, who reviewed earlier version of this work.