adversarial

April 9, 2022

via electronic submission

Jerry Perullo
Former Chief Information Security Officer **ICE/NYSE**
Former Chairman of the Board, **FS-ISAC**
Founder, **Adversarial Risk Management**
Professor of the Practice, Cybersecurity **Georgia Institute of Technology**

Securities and Exchange Commission
100 F Street, NE
Washington, DC 20549-1090

Re: **S7-04-22 - Cybersecurity Risk Management for Investment Advisers, Registered Investment Companies, and Business Development Companies.**

Ladies and Gentlemen,

I write to you as the recently retired Chief Information Security Officer of IntercontinentalExchange (NYSE:ICE), parent company of the New York Stock Exchange and a company uniquely positioned at the center of the global capital market ecosystem. While most of my experiences will be directly relevant to regulatory proposals affecting public companies and critical economic infrastructure, I see a benefit in harmonizing vernacular and parallel approaches across cybersecurity rules promulgated by the Commission. Further, as ICE CISO I also served as Chairman of the Board of the Financial Services Information Sharing and Analysis Center (FS-ISAC), which connected me closely with financial services firms of all sizes dealing with cybersecurity and resiliency. To that end, while I reserve my most substantive feedback for public company proposals, I am providing structural feedback that will be relevant to all cybersecurity rulemaking from the Commission here on **S7-04-22 - Cybersecurity Risk Management for Investment Advisers, Registered Investment Companies, and Business Development Companies.**

The Commission is right to perceive growing cybersecurity threats to advisers and funds and a need for action to limit risk to clients, investors, and the public more broadly. Further, it is right to look for best practices and expect a base level of security among overseen entities. Care should be exercised, however, to avoid repeating some of the mistakes made in drafting the first waves of cybersecurity regulation over financial market infrastructure. An over-emphasis on “policies and procedures” has historically driven efforts to quickly implement idealistic documentation that is poorly understood and incompletely adopted. While this approach facilitates audit and enforcement based on non-compliance with internal documentation, rulemaking should aspire to uplift cybersecurity culture, practices, and controls first. This more holistic approach is better captured by the term *cybersecurity program.* I thus recommend the Commission consider substituting the term “policies and procedures” with “a program” or “programs” as appropriate. A cybersecurity program should be defined to include documented cybersecurity governance,

strategy, policies, and controls". The following sentence from section **C. Overview of Rule Proposal** from the **Introduction**.demonstrates an example:

> Under the proposed rules, such an adviser's or fund's cybersecurity ~~policies and procedures~~ *program* generally should be tailored based on its business operations, including its complexity, and attendant cybersecurity risks.

This change enables the valuable consideration of business operations to be more practically implemented, as that practice is more likely to manifest in vital strategy and risk analysis exercises and documents created before and outside policy and procedure documents. While it is understood that the term "policies and procedures" permeates prior rulemaking outside cyber and is thus valuable to demonstrate continuity and consistency, it should be limited to explanatory material demonstrating how prior rules are being applied to the cyber domain. In the majority of detailed content, however, it is far more informative and less likely to drive negative behavior if a broader cyber *program* is discussed.

1. Should we exempt certain types of advisers or funds from these proposed cybersecurity risk management rules? If so, which ones, and why? For example, is there a subset of funds or advisers with operations so limited or staffs so small that the adoption of cybersecurity risk management programs is not beneficial?

> There is not a need for exemption thanks to the consideration of tailoring a program to business operations and complexity. The Commission was wise to integrate these factors, as they not only relieve the extensiveness of a cyber program expected of smaller funds or advisors, but more importantly empower funds and advisors of all sizes to conduct the intelligence analyses and business studies that allow them to focus resources on specific threats and scenarios that are targeting them - versus being forced into a blind one-size-fits-all security approach.

2. Should we scale the proposed requirements based on the size of the adviser or fund? If so, which of the elements described below should not be required for smaller advisers or funds? How would we define such smaller advisers or funds? For example, should we define such advisers and funds based on the thresholds that the Commission uses for purposes of the Regulatory Flexibility Act? Would using different thresholds based on assets under management, such as $150 million or $200 million, be appropriate? Would another threshold be more suitable, such as one based on an adviser's or fund's limited operations, staffing, revenues or management?

> **Section A. Risk Assessment** adequately addresses this concern. It is important that during examination and enforcement it is recognized that the thoroughness of a Risk Assessment and impact of its conclusions should scale with the size of an entity.

3. Are the proposed elements of the cybersecurity policies and procedures appropriate? Should we modify or delete any of the proposed elements? Why or why not? For example, should advisers and funds be required, as proposed, to conduct a risk assessment as part of their

cybersecurity policies and procedures? Should we require that a risk assessment include specific components (*e.g.,* identification and documentation of vulnerabilities and threats, identification of the business effect of threats and likelihood of incidents occurring, identification and prioritization of responses), or require written documentation for risk assessments? Should the rules require policies and procedures related to user security and access, as well as information protection?

The proposed elements of cybersecurity *programs* are a good starting point, but will benefit from important refinements to drive positive behavior and avoid stopping at paper compliance.

A. Risk, Assessment

It is worth drawing the distinction between *threat objectives* and *risks*, both of which are important to sound Risk Assessment. Threat objectives organize threats by their motivation, such as extortion, sabotage, fraud, or data theft. Risks, on the other hand, relate to conditions within an environment that may allow a threat to materialize. *Extortion*, for example, is a threat objective while susceptibility to phishing is a risk.

Section (i) places an unreasonable emphasis on inventory tasks. While it is understandable to imagine that someone outside cybersecurity might expect an inventory of assets as one might see in physical security or financial asset management, it is the wrong mindset with which to approach cyber. As uncomfortable as it makes the uninitiated, cyber assets appear and disappear constantly in modern computing environments, with prevalent technologies such as cloud computing, serverless architecture, and containerization driving ephemeral computing that renders classical inventorying concepts obsolete. Security concepts such as zero trust acknowledge a dynamic computing environment and surround it with infrastructure and cloud-level controls such as segmentation, isolation, and identity-based entitlement models. 2015's **Regulation SCI** made the error of over-indexing on classification of systems, leading much examination and enforcement time to be wrapped up in trying to even define what a system is in an age of virtualization and chasing spreadsheets of constantly-rotating server names. This preoccupation with inventory has left little time for actually discussing the security controls associated with systems and none for the far more important matters of evaluating threats and testing actual scenarios. It would be unwise to perpetuate this error and have yet another body of rules that will require overhaul in the near future to gain efficacy.

The conceptual purpose of this section in the proposed rule can be retained by pivoting away from inventorying *components* and toward inventorying *threats* such as::

"Categorize and prioritize cyber threats based on analysis of threat intelligence, the business environment, critical services and data therein, and the resulting potential likelihood and impact of realistic cybersecurity scenarios."

Further within section A, "...require written documentation of any risk assessment" reads as "if you happen to perform any risk assessments, then you must document them in writing." This can have the unintended consequence of discouraging risk assessment activity. The goals of this section would be better served by "Periodic written assessment of the threats facing the firm should be performed no less frequently than annually". *Threats* are appropriate for this sort of periodic strategic assessment, while the proposed rules as written are appropriate to then go on and require assessment, categorization, and prioritization of *risks* separately, which should be a continuous process.

The wording of the proposal around third-party risk does well to identify the value of screening. Specifically, beginning with an analysis of which providers "receive, maintain, or process … information… or … access" is the right approach to avoid such a large volume of assessments that critical vendors are rushed through a tickbox evaluation process alongside less important ones. A lack of emphasis on initial screening and understanding the risk posed by a vendor before beginning analysis is plaguing the banking sector today and it will be critical to not repeat that mistake.

In section **B - USER SECURITY AND ACCESS** there is an important opportunity for improvement in section (2). While multifactor authentication is rightfully lauded for the critical protections it has brought against credential theft, those benefits are not a result of the "combination of two or more credentials". Rather, it was the coincidental introduction of "one-time passwords" such as digital key fobs and authenticator apps that uplifted security so substantially. In reality such dynamic credentials could be the sole single factor and retain the majority of their benefit, while "multiple factors" would be useless if none of them were dynamic. This section and subsequent discussion should be improved via "...implementing authentication measures that require users to present a *dynamic or one-time credential, such as that implemented via time-based one-time password (TOTP) applications, push-based authenticators, smartcards, universal 2nd factor (U2F), or their successor technologies* for access verification."

Later in section B, endpoint protection should be updated from "...inspects all files…" to "inspects all activity" to avoid a common mistake in focusing on at-rest artifacts when fileless or post-file malware is a more relevant threat.

Section **C - INFORMATION PROTECTION** - contains the only mention of an absolutely critical area - testing. An examination of global financial market infrastructure regulation would find testing to be a crucial component of new regulatory frameworks, no doubt inspired by CPMI IOSCO's comprehensive studies and amplification of intelligence-led testing requirements pioneered by the Bank of England. While many if not most global financial infrastructure regulators including our own Commodity Futures Trading Commission (CFTC) were heavily inspired by these global standards in drafting cybersecurity rules, the SEC was notably divergent in drafting Regulation SCI. The result is a significant blind spot that should be avoided in additional rulemaking. Rather than just noting that a program *could* include penetration tests, this area or a dedicated section should specify that firms should deploy Attack Surface Management tools, Bug

Bounty Programs, and/or Red Team testing against specific scenarios identified via threat objective analysis to be relevant to the firm and use findings from that activity to drive risk identification and remediation prioritization. In practice this form of continuous assessment is the only practice that has driven targeted meaningful program improvement, and by beginning with threat objective assessment this approach would accommodate smaller firms who may determine that they are at low risk of targeting and thus have few adversaries to emulate in testing.

Section **D - THREAT AND VULNERABILITY MANAGEMENT** could do well to capture the testing discussion from section C, and could be broadened in scope by being retitled to "**TESTING, RISK IDENTIFICATION, AND REMEDIATION**". In fact the term "vulnerability" has too often been associated with a specific class of automated "vulnerability scanners" searching for known software flaws. This practice and the associated "patch mania" overlooks the fact that security defects in packaged software account for only a subset of risks, all of which should be identified, prioritized, and remediated in parallel. Specifically configuration errors, default or captured credentials, errors in internally-developed software, and access control mistakes contribute equally or more than "patchable vulnerabilities" and all of them should be surfaced via testing and remediated timely.

4. Should there be additional or more specific requirements for who would implement an adviser's or fund's cybersecurity program? For example, should we require an adviser or fund to specify an individual, such as a chief information security officer, or group of individuals as responsible for implementing the program or parts thereof? Why or why not? If so, should such an individual or group of individuals be required to have certain qualifications or experience related to cybersecurity, and if so, what type of qualifications or experience should be required?

It can be more beneficial - especially for smaller firms - to establish a Cyber Governance *committee* with accountability for overseeing the cyber program than a single individual. In addition to providing responsibility this practice also encourages many traditionally outside cyber such as Chief Financial Officers or Counsels General to participate in setting the program mission and regularly review the results via risk assessments and incident reports. This practice should not discourage the establishment of a Chief Information Security Officer, but to recognize the spectrum of firm sizes and types under regulation a committee of existing executives can prove more universally attainable than a single new role, while simultaneously ensuring broad awareness and sponsorship of security.

13. Should we require that advisers and funds respond to cybersecurity incidents within a specific timeframe? If so, what would be an appropriate timeframe?

Across my extensive experience with global financial cybersecurity regulation I do not recall seeing requirements around response time. Regulation has usually focused on notification requirements and recovery time objectives (in the case of critical infrastructure). Complications with response time include the fact that most impactful

incidents suffer from significant lapses or lack of *detection*, which is a prerequisite for response. Further, there is little to no agreement or consistency in what response measures are appropriate, and this requirement could drive programs into codifying dangerous response practices to "stop the timer" such as prematurely disconnecting networks and systems before proper analysis is conducted.

Responding broadly for questions 14-19, it should be noted that third-party risk management is exceedingly difficult to accomplish with certainty. Credit should thus be given to strategies that attempt to mitigate risk from specific providers by limiting the access they have, identifying secondary providers and exercising activation plans, and surrounding third-party software or products with technical security controls and monitoring. Where reliance cannot be avoided, third-party risk management expectations should focus on requiring the firm to understand and articulate specific risk scenarios of concern and identifying controls and practices at firms that are germane to those specific threats. This practice of "intelligence-driven examination" is far more efficient and effective than trying to rate a one-size-fits-all "security barometer" for a given supplier. Firms conducting third-party risk management must be required to understand the services provided by a third-party and articulate the threat objectives posed by the relationship first and foremost to avoid the mistakes made in previous regulation.

Responding broadly for questions 20-25, the goals of annual reporting are sound but the implementation as written is far too prescriptive. In fact there should not be a specific new standalone report prepared just for the sake of compliance with this rule where the goals of it are met (or exceeded) by a suite of existing documented practices. A mature program will operate a regular review of threat objectives and program mission, use those conclusions to continually test controls, and use those results to drive remediation priorities. The artifacts from this practice will exceed the goals of this section, and it would be a diversion of resource to then repackage this work in a new report. This section should be reworded to describe the type of activity that is expected to be documented at least annually, and avoid mentioning prescribing "a report", leaving regulated entities to satisfy this by any combination of healthy practices.

E. FUND BOARD OVERSIGHT (questions 26-32)

The wording of this section amplifies the earlier issues noted with overuse of "policies and procedures". In reality, reviewing a cybersecurity *program* and strategy is far more appropriate for a Board, where as-written many firms will be driven to actually place the lengthy and specific policies promulgated on staff in front of Directors. As I pivot my career into Board Directorship and reflect on my experiences in front of our numerous subsidiary and parent Boards of Directors over the years, it is clear that actual policies are rife with technical detail germane to various job responsibilities, such as system administrator service account naming convention, software engineer input validation requirements, or acceptable random number generator algorithms for use in cryptography. It is not only unrealistic, but disingenuous to expect a Board to review or approve this sort of material. While governance is indeed needed over policy, it is appropriate for this to be performed by company management (often by the Cyber

Governance committee mentioned earlier). Board-level approval is appropriate for a cybersecurity mission and high-level priorities, best captured in a strategy document that summarizes the greater *program* and requires far less frequent updates than policies or procedures. Revising earlier sections to distinguish where it is a cybersecurity *program* that needs definition will set up this section for refinement. Separately, it is appropriate to expect periodic updates on the program *operation* - including critical risks and high-severity incidents - to be communicated to the Board (and often via a sub-committee).

4. RECORDKEEPING (questions 33-34)

The revisions to recordkeeping rules proposed here are sound, and wisely index on logs, data, and other artifacts related to an incident. This avoids needlessly dictating preservation of masses of irrelevant information but instead empowers regulated entities to determine the specific scope of data supporting an investigation for preservation.

B. Reporting of Significant Cybersecurity Incidents to the Commission (questions 35-44).

As written the proposal does a good job of defining incidents and severity to enable reasonable performance of notification. Requiring explicit mention of this or any other regulatory authority in cybersecurity policies and procedures, however, is a common error. In practice it is unwise to expect cybersecurity practitioners to keep up to date with regulation and compliance requirements day to day. Rather, cybersecurity policies and procedures should define criteria for incident escalation to legal or compliance staff that will ensure any incidents that are candidates for notification can be evaluated by the appropriate staff. Legal or compliance staff, in turn, can be required to maintain policies or procedures that recognize binding regulation.

C. Disclosure of Cybersecurity Risks and Incidents (questions 48-54)

While I will leave it to others to debate how widely and quickly incidents should be disclosed and to whom, there is significant jeopardy in the proposed rules about the disclosure of *risks*. Risks represent the *potential* for an incident to occur, and by definition spell out the condition which could lead to compromise. The disclosure of a critical risk with any modicum of specificity enough to make it useful to the reader would also directly increase the likelihood of it being exploited. There is little to no oversight benefit in seeing explicit risks, and this process could dramatically increase the attack surface of the Commission in addition to the regulated entity. Disclosure of risk assessment *activity* could prove appropriate, but actually communicating the state of risk at any time brings significant danger with little value and should be avoided outright. To be more in line with the Risk Factors section of a public company annual report, it may be appropriate to describe cybersecurity *threats* to the advisor or fund more broadly and

discuss the wider landscape around adversarial activity. It is important for the Commission to draw this distinction so there is no chance of firms mistakenly concluding they are expected to divulge their tactical weaknesses.

Thank you for this opportunity to publicly comment on the proposed rulemaking. After a career on the receiving end of a global swath of burgeoning cybersecurity regulation, I'm eager to transparently share lessons learned and make meaningful improvements in cybersecurity across the financial services sector and beyond.

Sincerely,

Jerry Perullo
Former Chief Information Security Officer **ICE/NYSE**
Former Chairman of the Board, **FS-ISAC**
Founder, **Adversarial Risk Management**
Professor of the Practice, Cybersecurity **Georgia Institute of Technology**
https://www.linkedin.com/in/perullo/