May 4, 2022

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

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

Re: **S7-09-22 - Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure.**

Ladies and Gentlemen,

I write to you as an independent cybersecurity expert and the recently retired Chief Information Security Officer of IntercontinentalExchange (NYSE:ICE), parent company of the New York Stock Exchange and owner/operator of multiple public equity markets and additional critical economic infrastructure entities inside and outside of the United States. During my 25+ year career I enjoyed the unique opportunity to liaise with international regulatory agencies including the Securities and Exchange Commission, with whom I was able to engage separately as an officer of a publicly traded US parent company and as the Chief Information Security Officer of two SEC-regulated critical infrastructure entities. Across my experience I've been able to take part in the broader regulatory lifecycle from rulemaking to enforcement across not only public equity markets focused on disclosure and shareholder reporting, but also core infrastructure focused primarily on mitigating cybersecurity impact. This experience sat at a unique juncture to observe what regulatory actions benefited not only public shareholders, but also citizens dependent on a sound, reliable, and performant critical infrastructure.

While the cyber regulatory frameworks over public companies and critical infrastructure are distinct, there have been significant lessons learned in critical infrastructure that should inform public company regulatory decision-making.

First, it is worth noting that critical infrastructure regulation shares a focus on incident disclosure alongside public company oversight. This captured the public eye during the Colonial Pipeline ransomware incident of May 2021 and subsequent Cyber Incident Reporting for Critical Infrastructure Act ratified in March of 2022. While the Act closed gaps in reporting for critical infrastructure in sectors lacking strong cybersecurity oversight, reporting requirements in many other sectors including financial services and energy have been enforced for many years. Those of us with history under "cyber-forward regulators" watched CISA's foray into critical infrastructure reporting with interest, as it took a fresh look at requirements similar to those

we've operated under for many years. While existing requirements governing critical infrastructure and other regulated sectors require notification of primary regulators, the new Act focuses on providing notification to DHS CISA. In proposal S7-09-22, the Commission goes a step further via notification to public shareholders. All of these efforts and their international counterparts wrestle with the same questions; What defines an incident? What defines materiality? And finally, when is it reasonable to expect notification after these criteria are met?

The Commission has done well to focus on *material* cybersecurity incidents. Other authorities have flirted with the notion of incidents that merely present the *potential* of jeopardy, panicking industries recognizing a *potentially* limitless expanse of reportable events. Materiality has been core to the Commission's remit since the Securities Act of 1933, and it is wise to extend this notion - which preserves signal-to-noise ratio for the investing public - to cybersecurity. When it comes to detailed requirements around reporting - namely timing, escalation, and the concept of aggregating low-severity alerts, I will supply direct feedback in line with the Commission's request for comment. First, however, it is essential to provide situational awareness about the cybersecurity incident response lifecycle.

Only a small subset of cybersecurity incidents begin with an observation of impact - such as a ransomware or denial of service attack taking systems offline. The majority of incidents come to light through an investigative process. Examples in this area include data theft, financial fraud, and resource hijacking. It is particularly notable that these sorts of incidents are the ones more likely underpinning the Commission's impression of underreporting, since more dramatic service-interrupting incidents are instantly observable from the outside.

Within this sphere of investigation-driven incidents, the majority begin with an unremarkable alert among many. The business of "triaging alerts" and sifting through the noise to determine which should get human attention is a core focus of cybersecurity innovation. These manual and automated processes "peel the onion" layer by layer, usually concluding that the observed activity is authorized, poorly documented, inaccurately categorized by nascent detection technology, or otherwise dismissed as a "false positive". Where investigations begin to confirm unauthorized activity or policy violations, however, they escalate in severity, bringing additional investigative resources to bear and expanding internal notification to a wider audience potentially including legal, compliance, and company management. Determining that damage occurred often happens only after this investigative process, at which point materiality can begin to be assessed. A mature incident response process will use escalation procedures to bring business leaders and legal or compliance professionals into the discussion at the appropriate phase, as they are the ones who can assess materiality.

It is important to have this model in mind and understand that incidents begin as tiny seeds, most of which die along the path of investigation. The majority of incidents that are ultimately deemed reportable will thus have had to matriculate through a process that adds information along the way until materiality can be properly assessed.

Separately, with regard to introductory material the Commission should be cautious in interpreting media reporting of an otherwise undisclosed incident as a lapse that should or would be corrected via the proposed rulemaking. There are numerous reasons for public reporting about a breach that do not correlate with materiality thresholds, ranging from the pure salaciousness of a rumor to false claims made by an adversary in pursuit of extortion. Directly or indirectly equating media reporting of an incident with materiality would put a new weapon in the hands of extortionists, since simply drumming up hype and a media story around a purported breach could trigger new public disclosure requirements and potentially move a stock price. In addition to confusing the investing public, this sequence of events would empower adversaries. Firms should be expected to apply well-documented criteria to incident classification and judgment of materiality with disclosure requirements tied to those conclusions, rather than the virality of a news story.

> 1. Would investors benefit from current reporting about material cybersecurity incidents on Form 8-K? Does the proposed Form 8-K disclosure requirement appropriately balance the informational needs of investors and the reporting burdens on registrants?

The proposal as written strikes a good balance. There will be times where an investigation is moderately jeopardized by reporting requirements, but thanks to materiality thresholds they will be limited to only the most dramatic and unavoidable edge-cases.

> 2. Would proposed Item 1.05 require an appropriate level of disclosure about a material cybersecurity incident? Would the proposed disclosures allow investors to understand the nature of the incident and its potential impact on the registrant, and make an informed investment decision? Should we modify or eliminate any of the specified disclosure items in proposed Item 1.05? Is there any additional information about a material cybersecurity incident that Item 1.05 should require?

The level of disclosure is appropriate. With respect to specifically asking if data or operations were impacted, it is worth noting that some incidents will have materiality due to one of those elements but not the other. For example, a ransomware attack may be reportable due to operational impact, but an assessment of data theft may conclude that portion is de minimis and immaterial. Likewise, a data theft incident may have a minor side effect of accidentally disabling individual immaterial systems in the process, and while the data loss may trigger materiality the operational impact would not. As written a registrant could specify that “no material data impact” or “no material operational impact” was observed in the two cases, respectively, but this ability should be considered and preserved if the language evolves.

> 3. Could any of the proposed Item 1.05 disclosures or the proposed timing of the disclosures have the unintentional effect of putting registrants at additional risk of future cybersecurity incidents? If so, how could we modify the proposal to avoid this effect? For example, should registrants instead provide some of the disclosures in proposed Item 1.05 in the registrant’s next periodic report? If so, which disclosures?

The “proposed timing of the disclosures” is better bundled into question 4 on the four business day requirement, but the treatment of the incident timeline among information to be disclosed warrants feedback. As written, requiring disclosure of “when the incident was discovered” is unlikely to consistently add to investor benefit while disproportionately carrying the potential to create confusion and generate questions that can not or should not be answered at the time of the 8-K filing. For example, given the incident lifecycle previously noted, it is possible the true discovery time of an incident via a raw alert was long ago, and only the discovery of additional details or activity later escalated into a candidate for materiality determination. In that case forcing disclosure of a distant date before the investigation was complete could drive speculation and second-guessing that self-perpetuates market impact. It would serve investors well for the Commission to require disclosure of “when the incident was declared material", providing discretion for registrants to include timeline detail where it is supportive and beneficial. It is right to permit this but not require this, as in some cases specific timelines would jeopardize investigations. The fact that reporting timelines are triggered by the assessment of materiality ensures that this timeline reporting leeway is not a concern for potential abuse.

Other information expected via the proposal is appropriate as written, but warrants support in case it is challenged during the rulemaking process. Specifically, the flexibility provided in “brief description of the nature and scope of the incident” allows the timing proposed to avoid exacerbating risk. Were the Commission to require details around vulnerabilities exploited or details results of investigative activity, this could quickly change and provide an adversary with useful detail to tailor their strategies.

> 4. We are proposing to require registrants to file an Item 1.05 Form 8-K within four business days after the registrant determines that it has experienced a material cybersecurity incident. Would the proposed four-business day filing deadline provide sufficient time for registrants to prepare the disclosures that would be required under proposed Item 1.05? Should we modify the timeframe in which a registrant must file a Form 8-K under proposed Item 1.05? If so, what timeframe would be more appropriate for making these disclosures?

The four-day notification period works due to the strict materiality test and because disclosed information is qualified via “to the extent the information is known at the time of the Form 8-K filing”. It is important those qualifications are not abandoned.

> 5. Should there be a different triggering event for the Item 1.05 disclosure, such as the registrant’s discovery that it has experienced a cybersecurity incident, even if the registrant has not yet been able to determine the materiality of the incident? If so, which information should be disclosed in Form 8-K based on a revised triggering event? Should we instead require disclosure only if the expected costs arising from a cybersecurity incident exceed a certain quantifiable threshold, e.g., a percentage of the company’s assets, equity, revenues or net income or alternatively a precise number? If so, what would be an appropriate threshold?

Mature incident response programs (including those driven to maturity via this Proposal) will design severity levels and escalation procedures so that the legal and compliance teams that ensure fulfillment of these obligations are brought into the incident response process when a potential trigger of these conditions is at hand. Any notions of triggering notification around technical timelines absent a materiality assessment are unsustainable, as they drive dangerous practices around trying to inject compliance personnel into the entire incident response process or expecting cybersecurity analysts to become regulatory enforcement experts. In practice the volume and sensitivity (imagine false positive alerts around routine suspicious employee behavior that are determined inaccurate) of low-level cyber alerts warrants limiting access to incident responders, and the criteria that incident responders adjudicate should be supportive but not directly reflective of codified regulation. Put another way, incident responders should know to flag the disclosure of data as a trigger for general escalation, at which point data protection officers, compliance personnel, and other legal professionals can assess the vast and constantly-evolving array of data privacy and other relevant rules such as those proposed. Any temptation to trigger reporting by incident discovery or - worse - the discovery of *potential* - destroys that balance and creates situations where the lines between incident responders and lawyers are dangerously blurred.

> 6. To what extent, if any, would the proposed Form 8-K incident reporting obligation create conflicts for a registrant with respect to other obligations of the registrant under federal or state law? How would any such conflicting obligations arise, and what mechanisms could the Commission use to ensure that registrants can comply with other laws and regulations while providing these timely disclosures to investors? What costs would registrants face in determining the extent of a potential conflict?

There will continue to exist regulatory frameworks across the globe that impose reporting requirements absent a materiality clause. In those cases it is possible that a registrant has to make a cyber-related disclosure through an independently-regulated foreign or US subsidiary despite not meeting the more prudent threshold of the Commission. Enforcement teams will be well-served to recognize this distinction and not view disclosures in other jurisdictions as breaches of duty under the proposed rule by default.

> 7. Should any rule provide that the Commission shall allow registrants to delay reporting of a cybersecurity incident where the Attorney General requests such a delay from the Commission based on the Attorney General's written determination that the delay is in the interest of national security?

This should be examined to determine if prevailing law already implies the same, which one would hope. If, however, that is not the case or too much is open to interpretation, this should be clearly spelled out.

> 8. We are proposing to include an instruction that "a registrant shall make a materiality determination regarding a cybersecurity incident as soon as reasonably practicable after discovery of the incident." Is this instruction sufficient to mitigate the risk of a registrant

> delaying a materiality determination? Should we consider further guidance regarding the timing of a materiality determination? Should we, for example, suggest examples of timeframes that would (or would not), in most circumstances, be considered prompt?

In light of the cybersecurity incident response lifecycle, materiality can not be assessed until an incident has escalated to a sufficient level to involve appropriate senior resources. While one might argue compliance with this statute by declaring it *impracticable* to assess materiality until an incident has matriculated through the escalation process sufficiently, it could drive counterproductive behavior if interpreted poorly as-written since a relatively insignificant alert can start the clock on "discovery". The issue is resolved via simplifying to "a registrant shall make a materiality determination regarding a cybersecurity incident as soon as reasonably practicable to do so".

> 9. Should certain registrants that would be within the scope of the proposed requirements, but that are subject to other cybersecurity-related regulations, or that would be included in the scope of the Commission's recently-proposed cybersecurity rules for advisers and funds, if adopted, be excluded from the proposed requirements? For example, should the proposed Form 8-K reporting requirements or the other disclosure requirements described in this release, as applicable, exclude business development companies ("BDCs"), or the publicly traded parent of an adviser?

On balance the answer here is no. Different regulations and rules have different stakeholders and missions to satisfy, and trying to view some as superseding others can have unintended consequences. For example, incident reporting proposals for Investment Advisors would only trigger materiality for a publicly traded parent a subset of the time, and in those cases the additional audience and standard filing types and locations under this proposal will be appropriate to receive treatment.

> 10. As described further below, we are proposing to define cybersecurity incident to include an unauthorized occurrence on or through a registrant's "information systems," which is proposed to include "information resources owned or used by the registrant." Would registrants be reasonably able to obtain information to make a materiality determination about cybersecurity incidents affecting information resources that are used but not owned by them? Would a safe harbor for information about cybersecurity incidents affecting information resources that are used but not owned by a registrant be appropriate? If so, why, and what would be the appropriate scope of a safe harbor? What alternative disclosure requirements would provide investors with information about cybersecurity incidents and risks that affect registrants via information systems owned by third parties?

In some cases registrants will be reasonably able to obtain information to make a materiality determination about third-party incidents. This will most commonly be because they observe an impact. In those cases the proposal as-written would clearly apply. In cases where a third-party incident cannot be perceived or yet perceived as material, it is not reasonable to expect

disclosure and the proposal as written would not require it. There is no need for change, as a third-party incident naturally begets a first-party incident when it becomes relevant.

> 11. We are proposing that registrants be required to file rather than permitted to furnish an Item 1.05 Form 8-K. Should we instead permit registrants to furnish an Item 1.05 Form 8-K, such that the Form 8-K would not be subject to liability under Section 18 of the Exchange Act unless the registrant specifically states that the information is to be considered "filed" or incorporates it by reference into a filing under the Securities Act or Exchange Act?
>
> 12. We note above a non-exclusive list of examples that would merit disclosure under Item 1.05 of Form 8-K covers some, but not all, types of material cybersecurity incidents. Are there additional examples we should address? Should we include a non-exclusive list of examples in Item 1.05 of Form 8-K?

The list of examples needs improvement, which is of particular relevance since registrants will be hungry for them. The first 3 examples are unfortunately vague and undifferentiated. For example the first instance simply notes "an unauthorized incident that has compromised the confidentiality, integrity, or availability of an information asset…" The next goes on to list an "...incident that has caused.. damage…" This is clearly a subset of the previous example. The third example has an unauthorized party accessing, altering, or stealing data, which is captured in the "confidentiality" compromise in the first vague example. Further, the Commision wisely prefaces the examples with "if determined by the registrant to be material", but this needs more emphasis since the most popular reason driving referral to this section will be in pursuit of judging materiality itself. Likewise, negative examples would be equally valuable as positive in determining reportability. Specifically, registrants would be well-served to hear treatment of common examples such as:

- The unauthorized transmission of confidential data by an employee before resigning. While in this case the unauthorized breach of confidentiality would constitute review, the small volume of information, relatively low sensitivity (eg business contacts versus material trade secrets), and/or ability of the firm to challenge the offender and gain an affidavit of deletion may disqualify it from being deemed material.
- The detection of a ransomware infection on a single or limited number of computers that was successfully contained before it could become widespread. In this case while the incident response efforts may have been significant, the lack of material operational impact and reasonable assurance of containment could drive the conclusion that reporting is not required.
- A ransomware infection disabling the ability of a remote office to conduct normal operations for a two week period. In this case the breadth of infection or time alone would not trigger reporting requirements, but the operational impact of the incident would if it were deemed likely to ultimately result in a material financial impact.
- A denial of service attack rendering a primary corporate web site unavailable for two hours during a business day. While a calculation of the direct revenue lost through the incident may be unlikely to compute a material dollar value, if a reasonable investor

might find the resulting information about targeting of the company, customer behavior in response to the attack, and the registrant's ability to respond to the attack persuasive in investment decision making then it could be deemed material for reporting requirements.

- The discovery of a critical vulnerability in an Internet-facing information system that could have been exploited to cause catastrophic damage. While the processes and operations that discovered this risk and oversaw remediation should be described per other requirements of this proposal, a risk does not constitute an incident unless it manifested as such, and there is no incident disclosure requirement.
- The discovery that a recent acquisition had suffered an undetected compromise years prior with undetermined extent.

13. Should we include Item 1.05 in the Exchange Act Rules 13a-11 and 15d-11 safe harbors from public and private claims under Exchange Act Section 10(b) and Rule 10b-5 for failure to timely file a Form 8-K, as proposed?

14. Should we include Item 1.05, as proposed, in the list of Form 8-K items where failure to timely file a Form 8-K will not result in the loss of a registrant's eligibility to file a registration statement on Form S-3 and Form SF-3?

15. Should we require registrants to disclose any material changes or updates to information that would be disclosed pursuant to proposed Item 1.05 of Form 8-K in the registrant's quarterly or annual report, as proposed? Are there instances, other than to correct inaccurate or materially misleading prior disclosures, when a registrant should be required to update its report on Form 8-K or file another Form 8-K instead of providing disclosure of material changes, additions, or updates in a subsequent Form 10-Q or Form 10-K?

This proposal is useful to provide updates, but a few small adjustments can avoid misinterpretation. First, the word "potential" should be avoided in the following example: "Any potential material future impacts on the registrant's operations and financial condition;". This example would be more specific if changed to "Any *likely* material future impacts…"

Second, "Any changes in the registrant's policies and procedures" is unnecessary and overly broad. As I'll articulate in separate feedback, the notion of policy and procedural documents is a bit misportrayed, and in fact firms are likely to deploy specific technical controls, repair shortcomings in control deployment, and/or deploy more control testing to improve prevention or detection in light of an incident. It would be dangerous to disclose those specific changes, leaving a high-level summary here redundant to the actions that would be described in response to the previous bullet on remediation.

16. Should we require a registrant to provide disclosure on Form 10-Q or Form 10-K when a series of previously undisclosed and individually immaterial cybersecurity incidents becomes material in the aggregate, as proposed? Alternatively, should we

require a registrant to provide disclosure in Form 8-K, rather than in a periodic report, as proposed, when a series of previously undisclosed and individually immaterial cybersecurity incidents becomes material in the aggregate?

The concept of aggregating lower-severity incidents to reach a reporting threshold does not resonate with the cybersecurity community, where the misguided conflation of volumetric events with urgency often drives confusion. While we frequently tout metrics around the raw number of attempted compromises or malicious network packets, in the age of automation and high-speed computing frequency and recurrence does not correlate with the likelihood of success or impact. A properly instrumented Cybersecurity Incident Response Procedure will elevate or create a new “parent incident” of sufficient severity to trigger notification and reporting thresholds when impact begins, and there is not value in separately litigating the math around whether the sum of low-severity incidents infers a condition warranting reporting. Put into example, a high severity Denial of Service attack doesn’t manifest as a large number of discrete low-impact incidents; rather it is a single major incident. In the event previous detections kicked off while the attack was ramping up and were deemed inconsequential, they are likely to be aggregated into a single incident long before materiality is judged. Similarly, individual phishing incidents that compromise users may open and close as medium severity incidents, but if there is such a rash of them that it demonstrates a concerted campaign, that campaign will constitute a high-severity incident in itself with materiality assessed when there is impact. The only potential validation of the aggregation concept may be fraud, if hundreds or thousands of low-value fraudulent transactions add up to a material economic value. In that case, however, there would be a parent incident tied to the root issue and thematic condition, which would then be assessed for materiality. The idea of rooting out individual smaller incidents for some sort of aggregation test doesn’t hold a practical application that would warrant implementation as rule.

17. Should we adopt Item 106(b) and (c) as proposed? Are there other aspects of a registrant’s cybersecurity policies and procedures or governance that should be required to be disclosed under Item 106, to the extent that a registrant has any policies and procedures or governance? Conversely, should we exclude any of the proposed Item 106 disclosure requirements?

The Commission should take care in deploying the term “policies and procedures" for the purposes of the proposal. While the term has seen frequent use before and since the Securities and Exchange Act of 1934, in recent times it is associated with specific detailed documents when applied within operational risk management and cybersecurity teams. Subsequent language in the proposal around disclosing risk management and cybersecurity *strategy* or *programs* more broadly is more reflective of the aim here, and should be favored throughout. Traditional regulatory goals citing “policies and procedures" were often associated with defending against internal abuse, ethical violations, abuse of privilege, and similar threats of opaqueness such as insider trading. Establishing firm policies to manifest Commission rules was a large step, followed by implementation, compliance, and audit *procedures* to ensure

employees obeyed the rules. Adversarial Risk Management - and specifically cybersecurity - introduces a threat element outside the firms' control, and assurance programs are thus much broader than policy and procedure. In practice what a diligent third party (be it an Investor, Regulator, or even Board member) should aspire to observe is deliberate identification of cybersecurity threats and a resulting cybersecurity program, most often codified in a *strategy* document. That strategy should set the course for cyber risk management activity including active testing against identified threats, the continuous identification of specific risks that would allow a threat to manifest, and the well-instrumented remediation of those risks with consistent repetition. Such activities are not codified via *policies* at most firms, where policies relate more to specific behavior requirements and prohibitions placed on employees. Likewise, *procedures* are highly technical in nature, containing screen shots and URLs of internal tools and specific mouse clicks expected of staff. Cybersecurity procedures are more akin to field manuals than high-level program descriptions. The Commission should specify *program* or *strategy* more commonly, as much will get lost in translation by the time resulting code promulgates through the long tail of registrant companies, and the wrong choice of words could drive companies to not only over-share sensitive technical documents and irrelevant employee policies, but also to reactively draft and ratify generic policy documents en masse that do not reflect a company-level strategic imperative.

With that semantic change favoring *strategy* or *program* for most occurrences of "policies and procedures" in Section D, the content sheds the jeopardy of misinterpretation and becomes effective.

> 18. Are the proposed definitions of the terms "cybersecurity incident," "cybersecurity threat," and "information systems," in Item 106(a) appropriate or should they be revised? Are there other terms used in the proposed amendments that we should define?

This statement can be substantially reduced to "*Cybersecurity incident* means unauthorized occurrence or policy violation involving information systems that jeopardizes confidentiality, integrity, or availability". The removals in the proposed version serve several important purposes. First, removing "registrant's" avoids an accidental exclusion of third-party systems. Second, including "policy violation" captures a large number of common security incidents ranging from employee misuse of data to violation of third-party contractual provisions, all of which would fortunately be further subjected to materiality provisions for the purposes of the proposal. Finally, removing "of… information systems" avoids the risk of inadvertently excluding cyber attacks with a physical threat objective such as those upon operational technology.

> 19. The proposed rule does not define "cybersecurity." We could define the term to mean, for example: "any action, step, or measure to detect, prevent, deter, mitigate, or address any cybersecurity threat or any potential cybersecurity threat." Would defining "cybersecurity" in proposed Item 106(a) be helpful? Why or why not? If defining this term would be helpful, is the definition provided above appropriate, or is there another definition that would better define "cybersecurity"?

If *cybersecurity* were to be defined, it should not be done in a circular manner by employing subordinate terms such as “cybersecurity threat”. That said, while we struggled for years with the term “cyber” as practitioners, it has converged over time into common parlance. Attempts to define it here would not be helpful.

> 20. Should we require the registrant to specify whether any cybersecurity assessor, consultant, auditor, or other service that it relies on is through an internal function or through an external third-party service provider? Would such a disclosure be useful for investors?

It could be valuable to specify that using third-parties does not exempt registrants from any requirements herein, but requiring disclosure of which functions are performed in-house versus externally has the potential to cast an air of inferiority on outsourced services and disadvantage the many registrants who do not have the resources to build and maintain cybersecurity teams directly. Core functions that should be expected for internal performance amount to governance, and those functions are treated separately in the proposal.

> 21. As proposed, a registrant that has not established any cybersecurity policies or procedures would not have to explicitly state that this is the case. If applicable, should a registrant have to explicitly state that it has not established any cybersecurity policies and procedures?

Again, switching terminology to cybersecurity *programs* or *strategies* offers help.

> 22. Are there concerns that certain disclosures required under Item 106 would have the potential effect of undermining a registrant’s cybersecurity defense efforts or have other potentially adverse effects by highlighting a registrant’s lack of policies and procedures related to cybersecurity? If so, how should we address these concerns while balancing investor need for a sufficient description of a registrant’s policies and procedures for purposes of their investment decisions?

A firm should indeed have to explicitly state if it has not established a cybersecurity *program* or *strategy* (not policies or procedures) with supporting rationale. The jeopardy of this disclosure inviting attack is self-policing, since a firm that would be susceptible to attack must have some level of cybersecurity program in place. Further, changing terminology lowers the bar for being able to affirm this statement while driving beneficial behavior. It would be a reasonable use of resources for smaller firms to draft a *strategy*, while forcing them to implement *policies* is likely to drive them into downloading reams of boilerplate policy that is not reflective of culture and practices.

> 23. Should we exempt certain categories of registrants from proposed Item 106, such as smaller reporting companies, emerging growth companies, or FPIs? If so, which ones and why? How would any exemption impact investor assessments and comparisons of

> the cybersecurity risks of registrants? Alternatively, should we provide for scaled disclosure requirements by any of these categories of registrants, and if so, how?

With proper terminology the proposal can adapt to different sizes and categories of registrants naturally.

> 24. Should we provide for delayed compliance or other transition provisions for proposed Item 106 for certain categories of registrants, such as smaller reporting companies, emerging growth companies, FPIs, or asset-backed securities issuers? Proposed Item 106(b), which would require companies to provide disclosures regarding existing policies and procedures for the identification and management of cybersecurity incidents, would be required in annual reports. Should the proposed Item 106(b) disclosures also be required in registration statements under the Securities Act and the Exchange Act?

No.

> 25. To what extent would disclosure under proposed Item 106 overlap with disclosure required under Item 407(h) of Regulation S-K ("Board leadership structure and role in oversight")with respect to board oversight of cybersecurity risks? To the extent there is significant overlap, should we expressly provide for the use of hyperlinks or cross-references in Item 106? Are there other approaches that would effectively decrease duplicative disclosure without being cumbersome for investors?

There is indeed substantial potential for overlapping or redundancy. That said, the solution does not have to be limited to *incorporation by reference*. Rather, consider focusing on internal governance under proposed Item 106 and leaving matters of the Board to Regulation S-K. In practice there is far more required of senior management when it comes to cybersecurity governance than the Board. Specifically, the ratification of true policies is appropriate for an internal Cybersecurity Governance ("CyberGov") committee, where senior managers such as Chief Financial Officers, Counsels General, or Chief Operating Officers can accurately debate the consequences of employee behavior requirements and potential improvements in security. Board-level oversight, on the other hand, would be misdirected to read through lengthy policies, no less hundreds of discrete procedures. I've commonly observed many of the elements being proposed for the Board in this section to manifest in proxy statements alongside Director qualifications, and I think it would be appropriate to move the following to section E. Disclosure Regarding the Board of Directors' Cybersecurity Expertise:
Disclosure regarding whether and/or how the Board:

- approves the registrant's cybersecurity strategy and priorities
- is regularly made aware of the registrant's cybersecurity risks
- receives timely notification of urgent cybersecurity incidents
- oversees management's implementation and governance of cybersecurity programs

Making this change would allow section D.2 to focus on management oversight, at which point the material therein would be appropriate.

> 26. Would proposed Item 407(j) disclosure provide information that investors would find useful? Should it be modified in any way?

Disclosure requirements around Boardroom cybersecurity expertise can be viewed through one of two lenses: as the admission of some culpability or weakness; or as the confirmation of a noble behavior. As written, the proposal is not clear about which one it is pursuing. It would seem more congruent with the rest of the proposal that the latter is favored, in which case disclosure would indeed provide information that investors would find useful after the language is tailored appropriately.

> 27. Should we require disclosure of the names of persons with cybersecurity expertise on the board of directors, as currently proposed in Item 407(j)(1)? Would a requirement to name such persons have the unintended effect of deterring persons with this expertise from serving on a board of directors?

Directors are already disclosed in proxy statements, generally alongside matrices showing their credentials and participation in various subcommittees. Cybersecurity expertise should be viewed as a credential or qualification, in which case Directors should not harbor any qualm about its assertion (née “disclosure”).

> 28. When a registrant does not have a person with cybersecurity expertise on its board of directors, should the registrant be required to state expressly that this is the case under proposed Item 407(j)(1)? As proposed, we would not require a registrant to make such an explicit statement.

A requirement around holistic discussion of the Board’s involvement in cybersecurity, including elements ported from section D.2, would support a requirement to disclose the lack of cybersecurity expertise in the Board. This approach would allow a hypothetical firm with minimal cybersecurity risk to identify a committee where cybersecurity is deliberated, the usage of a cyber consultant or Board advisor periodically, and resulting conclusions. In practice it is difficult to imagine a public company devoid of cyber risk, but prudent to structure requirements to accommodate such a firm. Likewise, any registrant with obvious cyber risk that used this process to dismiss the threat would risk the market consequences of the obvious folly.

> 29. Proposed Item 407(j) would require registrants to describe fully the nature of a board member’s expertise in cybersecurity without mandating specific disclosures. Is there particular information that we should instead require a registrant to disclose with respect to a board member’s expertise in cybersecurity?

As the recently retired Chief Information Security Officer of ICE and the New York Stock Exchange in active pursuit of public Board Director positions, I naturally would stand to benefit

from mandating a laundry list of specific cybersecurity expertise and experience. It must be acknowledged, however, that Directors will continue to be business leaders and experienced strategic decision makers that can add value across the full spectrum of Director responsibilities first, and experts in cybersecurity or any other specific subject matter second. Further, shifting the tone from encouraging firms to *highlight* cybersecurity experience as a positive versus *disclosing it* as an admission will drive an organic process of identifying relevant experience across existing and new Directors. Finally, requiring registrants to disclose Board-level cybersecurity education - particularly that provided in-house within the context of the registrant's program - should encourage firms to provide briefings and enable existing Directors to expand their expertise.

> 30. As proposed, Item 407(j)(1) includes a non-exclusive list of criteria that a company should consider in determining whether a director has expertise in cybersecurity. Are these factors for registrants to consider useful in determining cybersecurity expertise? Should the list be revised, eliminated, or supplemented?

The list should be refined to prize adversarial risk management experience over policy or academic theory. Criteria should drive Boards toward selecting Directors with practical experience identifying threats, assessing susceptibility, closing gaps, and continuously testing to ensure management is not overlooking potential risks.

> 31. Would the Item 407(j) disclosure requirements have the unintended effect of undermining a registrant's cybersecurity defense efforts or otherwise impose undue burdens on registrants? If so, how?

No. The current state of cybersecurity has adversaries targeting any company of which they can discern the name, and there is not much room for increasing attention once a company is operating on any public stage.

> 32. Should 407(j) disclosure of board expertise be required in an annual report and proxy or information statement, as proposed?

Focusing on the proxy will suit more registrants' existing practices and avoid duplication.

> 33. To what extent would disclosure under proposed Item 407(j) overlap with disclosure required under Item 401(e) of Regulation S-K with respect to the business experience of directors? Are there alternative approaches that would avoid duplicative disclosure without being cumbersome for investors?

See previous comments about how cybersecurity expertise should be viewed as another business experience. To that end, it would not be irrational to consider treatment in 401(e), particularly if overall cyber governance, including the identification of committees that oversee cybersecurity, is retained in 407(j).

> 34. As proposed, Item 407(j) does not include a definition of the term “expertise” in the context of cybersecurity? Should Item 407(j) define the term “expertise”? If so, how should we define the term?

The examples provided do an adequate job of explaining the term.

> 35. Should certain categories of registrants, such as smaller reporting companies, emerging growth companies, or FPIs, be excluded from the proposed Item 407(j) disclosure requirement? How would any exclusion affect the ability of investors to assess the cybersecurity risk of a registrant or compare such risk among registrants?

No. The value at risk for an investor is correlated with the value invested, and not the size of the registrant. As such, small firms not prioritizing cybersecurity risk may actually embody the most pressing cybersecurity risk to an investor.

> 36. Should we adopt the proposed Item 407(j)(2) safe harbor to clarify that a director identified as having expertise in cybersecurity would not have any increased level of liability under the federal securities laws as a result of such identification? Are there alternatives we should consider?

Yes - this proposal is supportive of the notion that cybersecurity expertise should be considered an asset and not a liability.

Thank you for the opportunity to comment on this important proposed rulemaking. The amount of time and expertise that the Commission invested into this effort is evident, and the work is a substantial improvement over previous attempts at cybersecurity regulation in other divisions. The proposed rules focus on the right outcomes while avoiding over-prescription, and I look forward to being an investor, Director, and member of the public under them.

Sincerely,

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