February 15, 2017

Charles F. Leonard, Chief Operating Officer
Cybernance Corporation
12600 Hill Country Blvd, Suite R275
Austin, TX 78738

Introductory Remarks

The agencies' ANPR is a thoughtful and thorough examination of the key areas that, managed well, can significantly reduce cyber risk. Both the spirit and the content of proposed rules should be applauded. We believe firmly in applying scrutiny to these types of control systems, and our belief is borne out in results that not only reduce risk, but also improve overall operations.

At the same time, we are wary of control systems that strive for both broad application and specific implementation. Breadth and specificity are inversely related, and systems that aspire to both will be wracked with tensions that cause undue pain on those subjected to their application.

Throughout the following responses we advocate consistently for a general approach to risk management. Guidance to the covered entity is helpful in focusing attention on key areas, and the ANPR document does a good job of delivering such guidance. But guidance taken too far becomes prescription, and we encourage the rule makers to be vigilant about observing the boundary between, which is not often clear. In our view, deference to the general is a safe bet, and can be hedged by an overarching rule that demands documentation and support for each firms' specific decisions.

About Cybernance

Based in Austin, TX, Cybernance is a software firm that offers a Cyber Risk Governance platform to help organizations assess, measure, report, and prioritize efforts to reduce cyber risk. Used by boards of directors, C-suites, and key leaders in audit, risk, and legal compliance (as well as security), the Cybernance platform is the first software-based internal control system for cyber risk.

Questions on Sector-critical Systems

Question 11 – *What factors should the agencies consider in a measure of interconnectedness resulting in a system being determined as critical to the financial sector, and how should such factors be weighted? Commenters are asked to provide quantitative as well as qualitative support and analysis for proposed alternative methodologies, thresholds and/or factors.*

The complexity of global supply chains is dazzling, and the interconnectivity of

the financial sector exemplifies that complexity. In order to navigate the complexity it would be helpful to have a map, so to speak, of the relationships that define the landscape. It is intuitively clear that within the sector, there exist key players whose involvement is "critical path" and whose failure would result in a major impediment to the functioning of the system.

Agencies should seek to audit, inventory, and map these relationships to obtain a schematic of the industry. Such a map would represent a first-order understanding of pathways, dependencies and critical failure points. Many of these may be intuitively obvious, but many may not. The only way to understand the territory is to embark on a rigorous exploration.

In the process of building the map, other data can be collected. A relationship diagram could be enhanced by adding layers to describe transaction metrics (volume, frequency, etc.), organization attributes (headcount, revenue, sector, geography), and a host of other data. Maintained properly, such a map could serve as the backbone of a "common operating picture" for the financial sector. This capability is a fundamental component of cyber risk management and a keystone of situational awareness.

Questions on Cyber Risk Governance

Question 13 – *How would a covered entity determine that it is managing cyber risk consistent with its stated risk appetite and tolerances? What other implementation challenges does managing cyber risk consistent with a covered entity's risk appetite and tolerances present?*

Helpful guidance on this topic is offered in NIST SP 800-39, Guide for Applying the Risk Management Framework. In essence, the Guide discusses the notion that an organization must employ an internal control system for managing cyber risk (in this context the framework is alternatively NIST 800-53 or the CSF). But regardless which standard or framework is employed, the organization should establish a method for assessing, selecting, assigning, and implementing relevant cyber risk controls. These assessments should be risk-based, and involve a process that assigns responsibility and accountability to an individual who owns the control point.

Contrary to prevailing attitudes, Cyber Risk Governance is not just a technology issue – it requires attention from leaders in Internal Audit, General Counsel, Risk Management, Human Resources and Procurement, as well as IT & Security. Control ownership can be segmented and assigned on these lines; indeed many internal control systems align fairly neatly with these responsibilities. Control ownership, when assigned this way, becomes a natural outgrowth of the owners' existing job functions and skills.

With this method of applying risk controls, a reporting structure is established that uses risk assessment as its baseline measurement. Control owners who are tasked with assessing risk can then be guided toward an appropriate implementation of each control based on a risk tolerance. Risk threshold and tolerance, then, become the *lingua franca* of the cyber risk governance framework, and can be communicated up and down the organizational hierarchy.

The emergence of Cyber Risk Governance technologies has made it easier to programmatically implement cyber risk programs. Software platforms allow leaders to assign control ownership, conduct risk assessments at the control level, and report to executives and boards using terminology that enables productive discussions around risk tolerance. (Disclosure – Cybernance sells a software solution that allows organizations to conduct such automated assessments and reporting.)

Question 14 – *What are the incremental costs and benefits of establishing the contemplated standards for the roles, responsibilities, and adequate cybersecurity expertise (or access to adequate cybersecurity expertise) of the board of directors? To what extent do covered entities already have governance structures in place that are broadly consistent with the proposed cyber risk governance standards?*

The Board of Directors provides, among other things, oversight of strategic risks and the appropriate management of those risks. Whether risk comes in a financial, operational, or cyber flavor, it is still risk and should be managed using established risk management methods. Technology is only one component of the cyber risk issue, and seeking a technologist to advise the board on cyber issues focuses attention too tightly on the network component, away from the whole problem. If anything, an expert should be sought who understands the emerging field of Cyber Risk Governance.

Control systems like the NIST Cybersecurity Framework or the FFIEC Cyber Assessment Tool provide broadly applicable guidance on relevant controls that align with commonly establish governance structures. Entities should begin the search for automated systems for implementing controls. The Cybernance Platform provides a software-based system for implementation and ongoing monitoring of both NIST and FFIEC controls.

Organizations who use these tools find that they are helpful in applying internal control systems for cyber risk to existing organizational structures. Again, cyber is more than a technology issue and should be dealt with by groups of leaders with expertise in audit, legal, risk, policy and personnel management.

Questions on Cyber Risk Management

Question 15 – *The agencies seek comment on the appropriateness of requiring covered entities to regularly report data on identified cyber risks and vulnerabilities directly to the CEO and board of directors and, if warranted, the frequency with which such reports should be made to various levels of management. What policies do covered entities currently follow in reporting material cyber risks and vulnerabilities to the CEO and board of directors?*

We are not a financial institution but can comment on behalf of our clients. Many are still challenged to bridge the gap between tactical threat management activities (often heavily rooted in technology) and strategic, board-level concerns (based on corporate policy and risk). Specifically, the challenge lies in applying consistent criteria to determine what constitutes "material". Any reporting on threats, vulnerabilities or incidents should necessarily be framed in this context, so that executives and boards can consider the report with some perspective.

Using the guidance of cyber risk control systems like the NIST CSF, our clients focus first on developing and documenting some risk-based definitions of what constitutes a material event. From those definitions, they can create procedures for escalation and reporting that align with risk tolerance and are presented in that context.

This has several impacts. First, it forces the organization to devote some energy to thoughtful creation of policy, which is then documented and more likely to survive the inevitable churn of business. Second, it establishes the bedrock principles that can be evolved into more advanced procedures like incident response, organizational recovery objectives, and crisis management plans. Third, it creates surface area in which leaders outside of technology (risk management and internal audit, for example) can engage and apply their professional expertise to an area that traditionally lives behind a thick wall.

Question 16 – *The agencies seek comment on requiring covered entities to organize themselves in a manner that is consistent with the contemplated enhanced standards for cyber risk management. Besides the approach outlined in the ANPR, what other approaches could ensure that entities are effectively monitoring, measuring, managing, and reporting on cyber risk?*

The NIST Cybersecurity Framework identifies 5 specific "Functions" of an effective cyber risk management program. Namely, an organization must have capabilities that allow them to "Identify, Protect, Detect, Respond, and Recover" from cyber incidents.

These Functions can be subdivided in a number of ways, one of which is the "10 Functional Domains." Those domains are listed below, and each contains specific control points that may be implemented to guard against cyber risk. They can be described in summary using terms that begin to approximate recognizable roles within the enterprise:

- **Risk Management –** Chief Risk Officer and risk management related functions
- **Asset, Change & Configuration Management –** CIO and IT-related functions; tracking technology and ensuring proper configurations throughout
- **Identity & Access Management –** Human Resources / Security; tracking personnel and privileges
- **Threat & Vulnerability Management –** CISO; tracking and taking action against high-impact cyber threats
- **Situational Awareness –** CISO; building functions around monitoring and awareness of security status
- **Information Sharing & Communications –** General Counsel; creating and implementing policy around how information is treated and shared, internally and externally
- **Incident Response / Business Continuity –** Operations; designing, building, and testing organizational resilience plans
- **External Dependency Management –** Procurement/Purchasing; ensuring the application of risk controls to external relationships (supply chain, etc.)
- **Workforce Management –** Human Resources; applying proper screening, testing, and oversight of workforce
- **Cybersecurity Program Management –** Executives; creating and implementing enterprise policy that defines the execution of the previous 9 domains

Organizations who implement a framework like this (disclosure: our customers do, using our software platform) find that they achieve immediate results in terms of breaking down barriers between the functions named above. The individual controls of each domain are heavily networked across the domains, resulted in a web of dependencies between each function. These dependencies serve as checks and balances, a sort of connective tissue that binds functions closer together and increases the likelihood that an organization will be resilient in the face of cyber threats.

Questions on Internal and External Dependency Management

Question 17 – *The agencies request comment on the comprehensiveness and*

effectiveness of the proposed standards for internal and external dependency management in achieving the agencies' objective of increasing the resilience of covered entities, third-party service providers to covered entities, and the financial sector.

The agencies have suggested creation of a strategy to guide internal/external dependency management that integrates with the covered entities' overall ERM strategy. One interpretation would be that a covered entity is required to implement a cyber risk governance framework that aligns internal and external requirements, with reporting structures that assure adherence to an established risk tolerance, and consistency throughout. Many of the specifics covered – creating advanced awareness of assets, dependencies, relationships, etc. – are in fact high-priority controls recommended in frameworks like the NIST CSF.

Specific to internal dependencies, the agencies suggest that CEs should work to implementing mitigating controls on the inherent risk of the organization. The FFIEC Cyber Assessment Tool (CAT) offers the beginnings of such a method. Emerging solutions in the Cyber Risk Governance offer software-based methods for conducting the CAT and identifying specific controls that should be implemented. (Disclosure – Cybernance sells a software solution that addresses this issue using an automated version of the FFIEC CAT.)

Specific to external dependencies, the agencies recommend that CEs should seek to build an advanced understanding of the relationships between external partners and the internal assets to which those partners have access or potential impact. Once this understanding is achieved, the CE should begin to implement control requirements for their partners that mitigate against any risk presented by the dependency.

Although the internal and external approaches will differ in terms of how they are implemented – through policy and mandate vs. through contracts and negotiations – they are of a fundamentally similar character. In essence, the requirements stipulate that a company must create inventories of key assets, relationships between those assets, and any dependencies. These inventories and their attributes should be enriched with measures that describe the criticality, sensitivity and risk of the asset (or relationship). With an understanding of dependencies and risks posed by each, the CE should begin to implement mitigating controls that target any risks that exceed the established risk tolerance.

In abstract, this requirement resembles the content and spirit of the NIST CSF. We would suggest that the agencies avoid requiring specific, prescriptive controls for such an effort. A better requirement would be to enforce the

application of a risk-based framework to the enterprise, and to specify that it must include an integration of both external and internal cyber risk concerns. This is reflected nicely in the ANPR document, and could be refined (only a bit) by suggesting use of something like the NIST CSF or FFIEC CAT.

Question 18 – *What challenges and burdens would covered entities encounter in maintaining an internal and external dependency management strategy consistent with that described by the agencies?*

The first barrier to achieving the proposed objectives is common in all organizations – financial or otherwise – who are trying to build effective cyber risk governance programs. Many people view cybersecurity as a technology problem, and instinctively relegate it to the realm of IT. The reality is that cyber risk governance results from collaboration among many business leaders: internal audit, general counsel, risk management, human resources, procurement, and IT/security. The organizational and management challenge of aligning all these functions along the same priorities – including cyber among their existing responsibilities – can be a daunting undertaking. We see this challenge play out to varying degrees in all of our customers.

More specifically, the proposed requirements to maintain integrated strategies for internal/external dependencies will encounter a similar version of the same problem. Identifying internal dependencies will require communication and collaboration between IT executives and the business units who depend on IT assets. They will need to consult with risk managers who maintain criteria and methods for assessing and rating risks. Compliance analysts who ensure specific rules are followed will need visibility into the process, and auditors will strive for assurance that all of this activity is properly controlled and safeguarded. All of this activity should be guided by a standard internal control framework (like NIST) in order to ensure that each stakeholder is aligned to the same set of outcomes.

Of course, taking inventory of assets and dependencies is one thing, and maintaining a real-time awareness of those inventories is entirely another. The technological and logistical challenges implied are staggering. To be sure, CEs should absolutely conduct rigorous audits and recordkeeping around these activities, but we recommend thoughtful consideration of the tradeoff between burden and benefit. To borrow from the spirit of this requirement: the risk control strategy should align with the organization's overall strategy. In the same way, the risk control requirements should be designed so that their proper implementation improves overall operational efficiency. Effective ERM not only reduces risk, it improves process.

We would recommend that the agencies stop short of prescribing specific rules about how an in/ex dependency management program *must* be implemented. Rather, the rules might describe the need for executive leadership, representation from key functions (audit, counsel, risk) with appropriate enterprise authority, and application of a standard for risk-based decision making to all relationships and dependencies that touch the organization. Achievement of these objectives could be demonstrated by internal or external audits of the organizational control systems for cyber risk.

The NIST CSF recently released proposed changes in a "Version 1.1" that featured a significantly expanded focus on "supply chain" – another way of identifying external dependencies. This would be an appropriate framework for assessing whether an organization has implemented the proper risk control structures.

Question 20 – *What other approaches could the agencies use to evaluate a covered entity's internal and external dependency management strategies? Please be specific as to each approach.*

Strategies will be easier to evaluate when they can be held against a standard. The NIST CSF allows for organizations to select specific risk controls from a wide variety of recognized authoritative sources in order to address broad areas of risk. For example a CE might choose to implement a blend of controls from NIST 800-53, COBIT, and ISO 27001 to address the NIST CSF "Identify" function, and a similar blend of controls from the 20-CSC and C2M2 to address the "Respond" function. In this way, CEs are given the latitude to implement controls based on their unique risk profile and capabilities, and the simultaneous ability to align those controls with a standard framework (the NIST CSF).

Software platforms in the Cyber Risk Governance space provide organizations with this capability – to assess controls at the granular level and report an enterprise view of the control structure relative to NIST. (Disclosure – Cybernance sells a software solution that helps organizations align their control systems with NIST standards.)

Question 21 – *How would the proposed standards for internal and external dependency management impact a covered entity's use of a third-party service provider?*

Currently, approaches to external dependency management or third-party risk management are highly variable from one organization to the next. As a result there is no standard and no way of benchmarking compliance, whether for the vendor or the customer. Purchasers of services have been left to create their

own standards (often an ad hoc mix of various controls, managed in a spreadsheet). Vendors must then spend countless hours complying with each customer's homegrown methods – all of which are fundamentally the same (but different). We suspect this frustration in part is behind the agencies' questions in this category.

Based on these frustrations – which we hear again and again – we expect that both sides would welcome the reasonable application of standards. An objective, mutually agreeable framework for managing these relationships would drastically reduce friction throughout the lifecycle of the engagement. Furthermore, the ability to objectively gauge vendors' cyber risk management capabilities would lead to the emergence of higher quality products and services, and an incentive for purchasers to preferentially adopt those over others. (Disclosure – Cybernance sells a software solution that allows organizations to apply NIST standards to their vendors in order to understand potential cyber risks in the supply chain.)

To be clear, this optimistic view of a standard for external dependency management assumes that such a standard would prioritize non-prescriptive guidelines. Using a risk-based approach like that of the NIST framework allows individual organizations to make well-informed decisions based on their own internal risk criteria, rather than forcing compliance with one-size-fits-all prescriptions.

Question 22 – *What additional issues should the agencies consider related to internal and external dependency management and the covered entities' use of third-party service providers? How should those issues be evaluated by the agencies? Please be specific.*

Some agencies have published documents that outline recommended language for procurement contracts – specifically related to critical IT or OT systems. An example is "Cybersecurity Procurement Language for Energy Delivery Systems" [https://energy.gov/sites/prod/files/2014/04/f15/CybersecProcurementLanguage-EnergyDeliverySystems_040714_fin.pdf] from the Energy Sector Control Systems Working Group.

Of course this document focuses specifically on the electricity delivery sector and contains language specific to that industry. But underneath that language laid the bedrock principles of secure supply chain management, which can be broadly applied to procurement of any critical system. This document is not prescriptive in nature, but rather acts as a guideline or a potential checklist for procurement and contract managers to ensure that they're focusing on the right things. A document in this spirit may be of great use in helping CEs address 3rd

party risk.

Questions on Incident Response, Cyber Resilience, and Situational Awareness

Question 23 – *How well do the proposed standards for incident response, cyber resilience, and situational awareness address the safety and soundness of individual financial institutions and potential systemic cyber risk to the financial sector, including with respect to the testing strategies and approaches? How could they could be improved?*

As with many of the other areas of focus in this ANPR, activities around cyber response, resilience and awareness of the organization are heavily interlinked with the cyber risk governance program. The prevailing standards (whether NIST CSF or another) each clearly identify the need to 1) prepare a plan, 2) respond quickly and 3) remain vigilant at all times. The proposed rules align perfectly with these priorities and thus are, in our view, valid and reasonable.

We believe that the existing standards do a very good job of prioritizing response and resilience, and so question the need for a new, higher standard. In particular, where the ANPR delves into specific identification of the types of records that should be stored securely we think this is overly prescriptive. Furthermore, specifying the method or location of storage for such documents is unlikely to apply reasonably to all CEs. As with other responses in this RFI, we feel the best approach is the one taken by the NIST CSF – non-prescriptive guidance based on a sound risk assessment methodology.

This is not to suggest that we disagree with the need to treat certain types of information with greater caution and apply more stringent controls. We understand the importance of drawing differences between public, sensitive, confidential and "eyes only" information. But it is important to recognize that individual actors and organizations within any sector are likely to use varying systems for naming, classifying, and storing data as it pertains to their needs. Rather than forcing a new system on the sector as a whole, we advocate requiring the use of standard control systems that address a generalized risk-based approach classifying and segmenting information. The specifics of such an approach should be left to the organization, but the agency should reasonably expect that systems are defined by clear processes, procedures and documentation of any activity or decisions made.

Question 25 – *How do covered entities currently evaluate their incident response and cyber resilience capabilities? What factors should the agencies consider essential in considering a covered entity's incident response and cyber response capabilities?*

We often use the controls specified by a NIST-based framework called "C2M2"

(Cybersecurity Capability Maturity Model) that contains 10 functional domains (see our answer to Question 16). One of these domains in particular, Incident Response & Business Continuity, contains an outline of the types of policies, processes, and procedures that an organization should develop that focus on response and resilience in the face of a cyber incident. Evaluating a CE against these controls using the NIST Maturity Tiers provides a very good generalized assessment of the organization's overall incident response capabilities.

In many cases, these controls depend on successful implementation of controls in other "domains" such as risk management or asset management. This reality echoes our sentiment (seen throughout this RFI) that effective cyber risk governance is a multi-disciplinary approach to enterprise risk. So in this case a full understanding of incident response and resilience would be enhanced by a full assessment of the entire internal control system. (Disclosure – Cybernance sells a software solution that conducts a full assessment of the internal control structure and prioritizes areas that need further attention, including incident response and resilience.)

Question 26 – *How do covered entities currently evaluate their situational awareness capabilities? What factors should the agencies consider essential in considering a covered entity's situational awareness capabilities?*

Similar to the answer for Question #25 (Incident Response / Operational Resilience), we rely on general controls in the "Situational Awareness" domain of the NIST-based framework "C2M2". As noted in the ANPR, activity around response, resilience and awareness should be mutually reinforcing, and indeed that is the nature of these standard frameworks. In particular, controls for situational awareness should focus on creating internal definitions and requirements, implementing monitoring and reporting that meets those requirements, and enabling communication channels that consistently reach the appropriate stakeholders.

Based on the discussion of integrating external/internal dependencies with an overall strategy, it appears that the agencies are in search of what could be called a "common operating picture" (COP) for each CEs operational risk exposure. The degree to which a CE is able to create and maintain this COP will have a significant impact on situational awareness in general. So we view situational awareness as a natural outgrowth of the core activities discussed in the sections on understanding internal and external dependencies.

Rather than specifying and prescribing a definition for situational awareness that must be met by each CE, we recommend that the requirement take a similar form to the others we've discussed – that of a general control. Instead of

creating requirements that mandate certain types of monitoring or frequency of reporting, create rules that require monitoring and reporting based on enterprise-defined risk levels.

An organization might be reasonably expected to 1) develop and maintain criteria for classifying sensitive information so that they can 2) assign risk designations that align with that level of sensitivity. Following that process, the organization should 3) create monitoring programs focused on remaining aware of threats that may target such information so that they can 4) apply protective measures and access controls to guard that information against threats deemed credible.

This general approach aligns with the NIST CSF, and allows for an adaptive approach that is much more likely to meet the needs of the population in aggregate without wasting resources.

Questions on Standards for Sector-Critical Systems of Covered Entities

Question 32 – *Should different RTOs be set for different types of operations and, if so, how? Should RTOs be expected to become more stringent over time as technology advances?*

RTOs should be based on the CEs detailed understanding of its customer's needs and expectations; they should not be prescribed. Rather, an organization should be required to have clearly defined RTOs, and to be able to demonstrate the decision-making process that led to the determination of each.

Such a non-prescriptive requirement recognizes the myriad differences from one organization to the next in terms of size, composition, services, customers, and established processes. It allows each organization to determine based on its own needs and capabilities how best to serve its function in the market.

Questions on Approach to Quantifying Cyber Risk Section

Question 34 – *What current tools and practices, if any, do covered entities use to assess the cyber risks that their activities, systems and operations pose to other entities within the financial sector, and to assess the cyber risks that other entities' activities, systems and operations pose to them? How is such risk currently identified, measured, and monitored?*

Our company (Cybernance) offers a software-as-a-service platform that assesses, measures, analyzes and reports the status of an organization's cyber risk controls. This Cyber Risk Governance platform is intended to bridge gaps between key stakeholders in managing cyber risk: the board of directors, C-suite, general counsel, internal audit, risk management, and of course IT and security

roles. Our customers have described Cybernance as "an internal control system for cyber risk management."

The platform uses standard frameworks for cyber risk management (like NIST) to identify the status of risk controls within the enterprise. With a baseline established, users are guided toward greater risk resilience through a prioritized list of control implementations. Because the platform is collaborative across many functional roles, it helps enterprise leaders understand where various functions should be working together in order to achieve greater cooperation, awareness and resilience.

Question 35 – *What other models, frameworks, or reference materials should the agencies review in considering how best to measure and monitor cyber risk?*

The agencies have identified NIST (both 800-53 and CSF) as a key framework, and our responses have revealed our strong conviction that NIST is a highly adaptive method. We recommend that the agencies pursue involvement in an emerging consortium of cyber risk leaders, the Cyber Analytics Institute (cyberai.org)

Question 36 – *What methodologies should the agencies consider for the purpose of measuring inherent and residual cyber risk quantitatively and qualitatively? What risk factors should agencies consider incorporating into the measurement of inherent risk? How should the risk factors be consistently measured and weighted?*

Mentioned above, the Cyber Analytics Institute exists for this very purpose. Using large quantities of aggregated data from internal cyber risk control assessments, CAI researchers focus on creating models for quantifying cyber risk and creating valuation metrics for investment in risk mitigating efforts. A key goal of this effort is to enable the application of models to insurance premium pricing, so that insurers can create financial incentives for behaviors that demonstrably reduce cyber risk. Given the agencies' interest in (and proximity to) this field, we recommend that they consider involvement with the CAI.

Questions on Considerations for Implementation of the Enhanced Standards

We believe in a non-prescriptive approach to risk management. Organizations, even those competing to serve the same customers, are inherently different and thus have different risk exposures. Although they should be required to maintain active and observable control systems for cyber risk, the specific nature of those risk controls should be left to the organization's leadership to decide.

Any regulation around cyber risk governance should provide guidance, much like this ANPR has, in specific functional areas (risk management, external/internal

dependencies, etc. We have discussed our approach to these functional areas at length.