1 New Orchard Road
Armonk, NY 10504-1722

December 13, 2013

Mr. Adam Sedgewick
Information Technology Laboratory
National Institute of Standards and Technology
100 Bureau Drive, Stop 8930
Gaithersburg, MD 20899-8930

VIA EMAIL: csfcomments@nist.gov

RE: Preliminary Cybersecurity Framework Comments
Docket No.: 130909789-3789-01

Dear Mr. Sedgewick:

IBM appreciates the opportunity to respond to the National Institute of Standards and Technology's (NIST) Request for Comments on the Preliminary Cybersecurity Framework ("Preliminary Framework" or "Framework"). The Framework development process under Executive Order 13636 has been a valuable exercise in collaboration among government and industry stakeholders. IBM applauds the initiative as a means to leverage industry expertise to assist organizations in identifying and managing cybersecurity risk in a structured and thoughtful way. To date, we are encouraged by the Framework effort and by the prospect that it will have a positive impact on the cybersecurity posture of the United States. We look forward to continuing to work with NIST and our industry partners on this effort and helping to refine future versions of the Framework.

The goal of the Framework is to provide guidance to critical infrastructure organizations to help them manage cybersecurity risk (*see* Preliminary Framework, lines 84-85). Staying true to its voluntary and flexible nature, the Framework does not dictate specific technologies, measures, or outcomes. Rather, it provides a common language for organizations to evaluate their cybersecurity posture and to identify and prioritize opportunities for improvement, recognizing that different entities and sectors will use the Framework in different ways (lines 96-98, 200-205). This risk-based, adaptable nature is critical for supporting market-driven innovation and avoiding a rigid, check-the-box compliance mentality that would incentivize minimal action and ultimately weaken security.

IBM believes the Preliminary Framework is a step in the right direction for helping critical infrastructure entities better protect themselves from cyber threats. As it is currently written, the Preliminary Framework will help guide organizations in building or improving their cybersecurity risk management processes, prioritizing risks that need attention, and optimizing cybersecurity expenditures. It can be used both by organizations with less mature cybersecurity capabilities to establish a strong risk management program, as well as by organizations with sophisticated processes as a tool for continued assessment. It also serves as a valuable resource and collection of informative standards, guidelines, and practices.

The Framework, however, needs clarification in certain key areas to ensure that it appropriately reflects its risk-based approach, is not misunderstood to require particular outcomes, and does not overreach in scope. We address these and other issues in our comments below.

Successful Elements of a Risk-Based Approach

IBM reviewed the Framework from the perspective of a company that both secures its own globally integrated enterprise (spanning 150 countries, more than 400,000 employees, 120,000 servers, and half-a-million networked devices), as well as provides security services and solutions to virtually every sector of business worldwide. Based on our experience with IT risk management both internally and externally, we believe that the Preliminary Framework is a reasonable representation of the elements that are required for a successful cybersecurity risk management program. For companies with established cybersecurity risk management programs, there is value in reviewing the Framework categories and the informative references at the subcategory level. For organizations with less advanced cybersecurity risk management programs or that are looking to create a cybersecurity risk management process from scratch, the Framework is a pragmatic model and method of implementation.

It is important to emphasize that the value of the Preliminary Framework is as a risk management program, not a static list of cybersecurity controls. Lists of specific controls often do not address the actual risks businesses face every day and cannot keep pace with the world's constant change. Today's businesses – which are increasingly dependent on technology for the delivery of products and services – are particularly subject to the risks associated with swift adoption of rapidly emerging IT-enabled business paradigms.

A risk management process – such as the one outlined in the Preliminary Framework – provides a dynamic mechanism for capturing, analyzing, and acting on information related to potential causes and impacts of the changes around us. As the knowledge and understanding of this change increases, the resultant cybersecurity risk that we are faced with will decrease. But we must also recognize that knowledge and understanding alone cannot remove all cybersecurity risk. Acceptance of a reasonable level of risk is an important function in managing risk. In this, the Preliminary Framework provides its most significant value: As a risk management program framework, it recognizes that total risk elimination is often impossible; instead, it instantiates a process of risk management that allows for identification, prioritization, and remediation of reasonable risk to organizations.

In keeping with this thought, the Preliminary Framework can be improved, over time, to better delineate best practices in the area of risk identification and risk measurement. For example, the Identify function emphasizes the need to inventory and prioritize business critical IT and information assets, but it is largely silent on how to identify risk to these assets within a business process context. Also, while the Respond function identifies the importance of incident response policies and processes, the Framework is again largely silent on the need to collect and analyze incident related data for the purpose of feeding actual losses into the risk management identification phase.

Along these lines, the Framework categories and subcategories should be clarified to reflect the Framework's risk-based approach. If taken out of context of the Framework's overall risk-conscious method, for example, a number of the subcategories could be read to suggest more of a blunt instrument approach that organizations should apply equally to all assets and data regardless of their importance or to all vulnerabilities regardless of their risk.[1] As discussed, however, most organizations will not be able to protect all data and assets successfully all of the time or eliminate all vulnerabilities. Organizations should therefore prioritize their efforts based on criticality and risk. While implicit throughout the Framework, this concept should be expressed more clearly in Appendix A.

It is also important to recognize that not every subcategory will be of similar value to each organization and that there are likely to be additional subcategories not presently included in the Framework that may be significant for some entities. Over time, as organizations sharpen their risk management programs and continue to learn from experience, they will help identify areas where the Framework itself can be improved. For this reason, it is important that the Framework recognize, as it currently does, that it is intended to complement, rather than replace, an organization's existing cybersecurity risk management process (lines 100-101). Indeed, companies should be encouraged to continue developing, updating, and employing their own cybersecurity risk management frameworks. As organizations fine-tune their efforts to their own needs and risks, their market-driven experience will benefit future versions of the Framework.

The Framework can thus be supplemented over time with a repository of institutional knowledge related to cybersecurity risk that can help organizations "jump-start" and/or optimize cybersecurity risk management activities. Until that time, organizations can turn to the Framework – with the assistance of trusted advisors like IBM – for guidance on how to drive the organizational focus, process, and management foundation necessary to manage prevalent risk. IBM Security Services has already released a maturity assessment based on the Preliminary Framework as written today and plans to formally release an assessment based on the final Framework in February 2014.

Adoption

Much of the discussion at the 5th Cybersecurity Framework Workshop focused on what it means for an organization to adopt the Framework. A number of participants expressed concerns that the meaning of "adoption" was too uncertain and suggested clarifying the definition in the Framework.

We agree that the definition of "adoption" should be clarified. The concept of adoption itself can be confusing in the context of the Framework because it implies a fixed set of rules and an outcome or performance-based metric that can be universally applied across all critical infrastructure entities. But such a metric would be in conflict with the very nature of the Framework as a flexible and adaptable tool. As described below, what we believe is really meant by "adoption" is the *utilization* of the Framework as a guide to identifying and managing

[1] *See, e.g.,* ID.RA-1 ("Asset vulnerabilities are identified and documented."); PR.DS-1 ("Data-at-rest is protected."); PR.DS-2 ("Data in motion is secured."); PR.DS-5 ("There is protection against data leaks."); PR.DS-6 ("Intellectual property is protected.); PR.PT-4 ("Communication networks are secured."); DE.CM-4 ("Malicious code is detected.")

cybersecurity risk. The "utilization" concept is consistent with the definition that NIST recently provided in its Update on the Development of the Cybersecurity Framework (Dec. 4, 2013) (discussed below).

Consistent with its risk-based approach, the Framework sensibly recognizes that organizations vary widely in their business models, resources, risk tolerance, approaches to risk management, and effects on security, national economic security, and national public health or safety (lines 180-182). As a result, the Framework allows for flexible implementation, noting that different entities will use it for different means (lines 203-205). Importantly, it does not direct particular actions to be taken or require certain outcomes to be achieved. Rather, the categories and subcategories listed in Appendix A are activities and outcomes to be *evaluated*, and the corresponding standards, guidelines, and practices are non-exhaustive example sets of common ways organizations may choose to address risk in those categories (lines 235-237). In other words, the Framework is "not a checklist of activities to perform" (line 208).

What the Framework is, instead, is a common language for expressing, understanding, and managing cybersecurity risk (lines 200-201). At its essence, it is a methodology – a tool – for conducting a self-assessment of cybersecurity risk (line 148). It helps organizations identify their current cybersecurity posture for known categories of risk, set goals for their target state for such categories, and identify and prioritize actions for reaching the target state (lines 96-98). It provides informational resources to help guide organizations in achieving their risk management goals, but each organization must choose what its goals are and how it reaches them. Even the risk management process itself will vary across organizations (lines 100-104).

Some of the confusion about adoption may result from the fact that the categories and subcategories are described with outcome-oriented language. For example, subcategories PR.DS-1 and RS.MI-2 simply state the outcomes: "Data-at-rest is protected," and "Incidents are eradicated." As with all the other subcategories, however, the intent of this language is not to state an outcome that must be achieved before an organization is considered to have adopted the Framework. Protecting data-at-rest, of course, can be a very complicated and challenging effort in which risk can be mitigated but may never be fully eliminated. And eradicating incidents is an ideal end-state that organizations should constantly strive for but may have difficulty reaching permanently. Other subcategories may be more realistic to achieve but of varying importance in the risk profiles of different organizations. Many entities will even have varying and constantly changing levels of maturity within their own organization, as they adapt to events such as mergers and acquisitions or the introduction of new products, services, or business ventures. The subcategories therefore should be read not as static outcome requirements, but as outcome-oriented goals that organizations can assess, prioritize based on their risk profile, and evaluate progress towards, while using the informative resources as examples of practices that may help them achieve progress.

Confusion about adoption may also result from misinterpretation of the Framework Profile and Framework Implementation Tiers. Similarly, these should be tools for organizations to set goals, evaluate their own progress, make informed decisions, and prioritize resources. They should not be seen as a method for calculating standard scores or metrics that organizations must reach to adopt the Framework. Each organization may have different priorities and may measure their profiles and progress in different ways. Depending on the organization, lower maturity levels on particular subcategories may be an appropriate risk-management strategy.[2]

To address the confusion and concern about adoption, we recommend that the Framework contain a clear statement defining what it means to adopt the Framework. In IBM's view, organizations should be considered to have "adopted" – or, more aptly, to be "using" – the Framework if they have incorporated the Framework's main *evaluation* concepts into their own cybersecurity risk management processes. To be more specific, an organization uses the Framework when it has incorporated into its own cybersecurity risk management program *processes* for identifying, assessing, prioritizing, and/or communicating its: (1) risk posture with respect to the Functions listed in the Framework Core; (2) current and target approaches to address those risks; and (3) cost-effective actions to reduce those risks considered in the context of the organization's broader risks and priorities.[3]

We also recommend that the Framework clarify a few related points. First, it should state that it is not intended to create a universal metric of adoption. To emphasize this point, it should clarify that there are no defined sets of risk mitigation actions that organizations must take, or outcomes that organizations must achieve, in order to adopt the Framework.

Second, the Framework should make clear that the categories and subcategories, while described in outcome-oriented language, are goals to be set and evaluated. The Framework should recognize that there will be a spectrum of progress towards those goals, and many of them – like "Data-at-rest is protected" – may never be fully achieved. It should emphasize that the point of the Framework is the *process* for evaluating these goals, and that organizations may vary considerably in their progress and approach based on their own business models and resources, among other factors.

Third, the Framework should make clear that the Framework Profiles and Tiers are tools for organizations to use internally to prioritize actions and measure their progress, and that they should not be read to require a particular outcome or score.

Voluntary Nature

A fundamental aspect of the Framework is its voluntary nature, which enables the Framework's flexible and risk-based approach. Any attempt to mandate the Framework would

[2] There is also ambiguity in the Framework about (1) whether the Tiers apply to each subcategory or to an organization's overall cybersecurity risk management maturity, and (2) whether there should be a standard methodology for measuring the Profiles and Tiers across organizations. Clarification of the first issue depends on part on the second. While promoting a common language is an important goal of the Framework, we believe it is too early in the Framework development process to describe a standard methodology for measuring Profiles or Tiers. Many entities today use different and valid methods for measuring cybersecurity maturity, and NIST should encourage the market to continue developing those methods before suggesting a particular method in the Framework.

[3] This is largely consistent with NIST's proposed definition in its December 4 update, with a few minor changes. *See* Update on the Development of the Cybersecurity Framework (Dec. 4, 2013) ("An organization adopts the framework when it uses the Cybersecurity Framework as a key part of its systematic process for identifying, assessing, prioritizing, and/or communicating: cybersecurity risks; current approaches and efforts to address those risks; steps needed to reduce cybersecurity risks as part of its management of the organization's broader risks and priorities.").

likely encounter efforts to impose certain outcome-based requirements. Such a prescriptive regime would result in weaker cybersecurity programs by encouraging firms to invest only in meeting specific items on a checklist that ultimately would not be able to keep pace with rapidly changing technology and threats. As the President's Council of Advisors on Science and Technology (PCAST) stated as its central conclusion in its November 2013 Report to the President, cybersecurity will not be achieved by a collection of static precautions or list-based mandates that encourage a "check-the-box" mentality and provide incentives for minimal compliance, but rather requires a set of processes for continuous improvement.

Standing by itself, the Framework represents the type of voluntary, continuous risk-management process that will help organizations build or improve on their existing cybersecurity programs in a flexible way that makes sense for individual business and risk profiles. If used as intended, it will support efforts by companies to prioritize risks that require attention and optimize cybersecurity expenditures, while avoiding the pitfalls and false sense of security of rigid or prescriptive compliance requirements.

The potential for counterproductive mandates, however, exists beyond the development of the Framework. Executive Order 13636 requires federal agencies to evaluate the Framework against current regulatory requirements for critical infrastructure and to propose new requirements where current requirements are deemed insufficient. The Executive Order also required a review of existing federal procurement requirements related to cybersecurity.

For the reasons already discussed, we strongly discourage any efforts to transform the voluntary Framework process into a mandated program. Mandating the Framework in federal agency procurements, for example, could have the effect of focusing acquisition and program management efforts on compliance-related metrics, leading to reduced flexibility for government contractors to employ risk-based and innovative efforts to improve cybersecurity. Any reference to the Framework in a regulatory, procurement, or other related context must not lose sight of what the Framework is: a language and structure for risk management processes that will be implemented differently in each organization, not a checklist of specific measures that must be taken or outcomes that must be achieved. For this reason, it is important for NIST to provide the clear guidance recommended above on the issues related to "adoption."

Standard of Care

Related to the issues of adoption and voluntariness is the legal question of whether the Framework will set a new standard of care in cybersecurity-related litigation or regulatory proceedings. Some observers have assumed that the Framework, even though voluntary and non-regulatory in nature, will nonetheless become a de facto standard of care that courts and regulators will look to for guidance on what constitutes reasonable cybersecurity measures.

For the same reasons that adoption of the Framework cannot be measured by the implementation of a particular control or the achievement of a specific outcome, the Framework should not be viewed as setting a new legal standard of care for cybersecurity liability or compliance. For example, a plaintiff in a civil action or a regulatory agency should not be able to point to the Framework itself as evidence that an organization should have implemented a particular control, achieved a stronger outcome, or set a higher priority for a particular category. The Framework does not dictate certain measures or outcomes and should not be misunderstood as doing so.

NIST, of course, cannot control how litigants and others will try to use the Framework in adversarial settings. But given the significant potential for misinterpretation of the Framework in such contexts, the Framework's intent should be clarified on this issue. To avoid such confusion, we suggest that the Framework include language stating that it is not intended to set a legal standard of care on any particular cybersecurity measure or outcome, or define what constitutes "reasonableness" in any particular context, as such a result would be inconsistent with its purpose of providing a flexible, risk-based approach that will be adapted and used by different organizations in different ways.

Privacy

IBM has a long history of privacy leadership – from our adoption of one of the world's first global privacy codes of conduct in the 1970s, to our status as the first company to be certified under the Asia Pacific Economic Cooperation Organization Cross Border Privacy Rules earlier this year. We agree that the Framework should address privacy; doing so will bolster public trust in, and support for, cybersecurity programs. We believe, however, that the Framework should not attempt to set forth a comprehensive approach for implementing a privacy program but rather focus on the *privacy impacts of cybersecurity actions.*

The privacy approach in the current draft is widely seen as too broad and beyond the appropriate scope for the Framework. A targeted approach is more suitable because not all cybersecurity actions have privacy impacts. Moreover, privacy considerations vary according to context, jurisdiction, and sector and are evolving rapidly. Different industries have differing approaches and future needs; they will hesitate to adopt a single approach to privacy, especially one marked by broad and open-ended direction, such as that proposed by Appendix B. Because a targeted approach would make the Framework more compatible with varying privacy regimes, it would make the Framework more attractive to organizations operating in multiple jurisdictions and more influential internationally. A more targeted privacy approach would also help avoid concerns about creating an unintended regime of de facto privacy regulation.

By focusing on privacy impacts of cybersecurity actions, the Framework can provide clearer and more actionable guidance that will be more readily adopted by organizations across a variety of industries. Such guidance (whether incorporated in Appendix A or listed separately) could appropriately address categories such as:

- Assessing the privacy implications of an organization's cybersecurity program;
- Adopting processes to support compliance of cybersecurity activities with applicable privacy laws;
- Adopting processes to identify and mitigate potential privacy impacts of monitoring, gathering, or sharing information for cybersecurity purposes; and
- Appropriately training individuals with cybersecurity-related privacy responsibilities.

* * *

IBM appreciates the opportunity to provide these comments on the Preliminary Framework and looks forward to further collaboration with NIST and our other partners on this effort.

Sincerely,

Brendan Hannigan
General Manager, IBM Security Systems

Kris Lovejoy

Kris Lovejoy
General Manager, IBM Security Services

Joanne L Martin

Joanne L. Martin
Vice President and IBM Chief Information Security Officer

Andrew Tannenbaum
IBM Cybersecurity Counsel