Policy Paper

April 2021

Quantifying Risk: Innovative Approaches to Cybersecurity

Adam Bobrow

G | M | F The German Marshall Fund of the United States
STRENGTHENING TRANSATLANTIC COOPERATION

Washington, DC Ankara Belgrade Berlin Brussels Bucharest Paris Warsaw

Summary

The Framework for Improving Critical Infrastructure Cybersecurity, published by the National Institute of Standards and Technology (NIST) in 2014, helped drive a heightened focus on risk-based approaches to addressing cybersecurity concerns. The framework's impact on the practice of cybersecurity has been universal but diffuse: while it put risk at the center of how organizations across the economy think and talk about cybersecurity, it did not provide the tools for cybersecurity professionals to measure and manage that risk. By any measure, the cyber ecosystem is in worse shape today than ever, and efforts—including the framework—to improve cybersecurity have not garnered the hoped-for results.

In virtually all other domains where critical assets or infrastructure are under threat, risk management involves rigorous, quantitative analysis to direct cost-effective investments in prevention or mitigation. Despite citing risk as the most important means to assess cybersecurity, quantitative risk analysis remains absent from most cybersecurity programs today. Neither the awareness of the problem nor the need for practical risk-based approaches to solve it is new. To cite just one example, a 2003 report on challenges in achieving trustworthy computing identified the need for the development of risk measurement for cyber risk within a decade. Instead, the situation today is worse than it was in 2003 when the experts met.

Making risk-based cybersecurity decision-making a reality requires the development of a measurement system that allows meaningful comparisons of risk among different organizations across industries. Even in the few organizations that make the effort to build metrics based on cyber risk, those apply only to the organization for which they were developed. One organization cannot measure its risk in a way that aligns with how its peers do or conduct any normalization against a generic baseline.

Some of the groundwork for developing a common set of cybersecurity metrics has already been done. For example, a Department of Homeland Security working group of insurance industry experts, convened in 2015, described the characteristics of the data that needs to be collected. The insurance industry has thus far failed to implement the requirements identified in that effort.

Changing the status quo to enable meaningful organizational cybersecurity decision-making requires that the U.S. federal government play the role of honest broker and facilitate the development of a more quantitative approach to cybersecurity, including by:

- Collecting broad-based data about past incidents and releasing anonymized data sets based on incident reports that the private sector can use to build the tools and help organizations build cybersecurity capacity,
- Developing actuarial models that project the impact and likelihood of future incidents in quantitative terms, and
- Facilitating the creation of metrics to enable concrete comparisons within and among a diversity of organizations.

This paper offers four recommendations to the U.S. federal government, to jump-start the effort at quantifying cybersecurity risk and making true risk-based analysis of cybersecurity a reality:

- The president should issue an executive order charging NIST with revising the Cybersecurity Framework in a way that focuses on quantifying cybersecurity risk and the secretary of commerce with developing initial quantitative cyber risk models from available data sources.
- Congress should create a Bureau of Cyber Statistics at the Commerce Department with a mandate to collect private-sector incident information.
- Congress should create a National Cyber Safety Board to investigate cyber incidents.
- The federal government should run an open innovation grand challenge to demonstrate how quantitative models can lead to better cybersecurity outcomes.

The Status Quo Is Unsustainable

In 2003, the Computing Research Association (CRA) convened a group "devoted to defining technical and social challenges in trustworthy computing." Its report, titled Four Grand Challenges in Trustworthy Computing, concluded that the status quo in computing at that time was unsustainable because, for all the benefits provided by the increasing ubiquity of computing devices, the security of those devices was almost universally inadequate.[1] The report identified increases in the number of reported incidents and the cost of downtime from an incident, and identified four goals to reverse the rising levels of cyber insecurity, including the goal of "developing quantitative risk management techniques."

Eighteen years later, the trends identified in the CRA report have worsened, the cybersecurity of organizations around the world is weaker than ever, and quantitative cyber risk management techniques remain elusive. Because security was an afterthought in the development of the globe-spanning networks that have emerged over the last 25 years, cybersecurity emerged as a reaction to cyber incidents rather than part of the architecture of the Internet. In the early days of networked computers, system administrators of victimized networks developed practices to prevent the recurrence of similar incidents. Over time, these practices were standardized and aggregated into controls: policies designed to prevent incidents or lessen their impact, and allow for response during, and recovery after, an incident. For example, requiring passwords to access private networks was one of many practices incorporated into controls governing access and authentication.

Initially there was a naive assumption that, if system administrators fully implemented all relevant controls, nearly complete security would follow. As the volume and sophistication of cyber threats grew, however, the financial cost of control implementation increased to such an extent that no organization could hope to effectively maintain a system with "all relevant controls" fully implemented. At the same time, it became clear that the nearly complete security promised was not possible given the insecure foundation of modern computing infrastructure. With ever more sensitive data residing on vulnerable networks and increasingly critical systems accessible over the public Internet, cybersecurity professionals began to look for a new way of thinking about the problem. Three trends emerged to refine the control-based management of cybersecurity:

- A greater focus on protecting critical infrastructure as a way of prioritizing resources;[2]
- The range of activities that controls covered expanded beyond prevention and detection of incidents to address response and recovery in a more serious way; and
- Discussing cybersecurity as a risk-based discipline with comparisons to other fields where risk management has had a significant impact.

These three shifts were captured in the 2014 Framework for Improving Critical Infrastructure Cybersecurity, developed by the National Institute of Standards and Technology (NIST) at the direction of the Obama White House.[3] Executive Order 13636 instructed NIST to develop a framework to provide "a prioritized, flexible, repeatable, performance-based, and cost-effective approach, including information security measures and controls, to help owners and operators of critical infrastructure identify, assess, and manage cyber risk."[4] The Cybersecurity Enhancements Act of 2014, which codified the terms of EO

1 Computing Research Association, Four Grand Challenges in Trustworthy Computing, Second in a Series of Conferences on Grand Research Challenges in Computer Science and Engineering, November 16–19, 2003.

2 Cybersecurity and Infrastructure Security Agency, Critical Infrastructure Sectors, U.S. Department of Homeland Security (DHS), October 21, 2020.

3 Executive Order No. 13636, Improving Critical Infrastructure Cybersecurity, February 12, 2013, §7(d).

4 Ibid., §7(b).

13636, charged NIST with continued maintenance of the framework.[5] NIST updated the framework by releasing a Version 1.1 in 2018, and it has continued to publish additional resources provided by the cybersecurity community.[6]

The framework adopted a rhetorical focus on risk as the frame of analysis for improving cybersecurity—identifying cybersecurity risk as a category of risk, along with financial and reputational risks, that can impact an organization's bottom line.[7] In doing so, the framework explicitly embraced a growing consensus that risk offered the best way to conceptualize the trade-off between the benefits of interconnected information systems and the inherent insecurity of those systems. The idea of risk, however, is overshadowed by the categorization of controls into a functional hierarchy. Rather than offering a thorough means to assess risk, the framework fits more comfortably alongside maturity models for control implementation.

While the framework was intended to improve the cybersecurity of critical infrastructure in the United States, that was far too ambitious an objective for the framework and the ecosystem remains largely insecure. Instead, the framework has helped reorient the discussion of cybersecurity toward risk in critical infrastructure sectors and across the economy. Cybersecurity professionals across a broad range of sectors now attempt to describe their programs as an effort to "determine the acceptable level of risk for achieving their organizational objectives" and "express this as their risk tolerance."[8] The framework, however, has not provided organizations with the tools to measure their risk, so the change in how cybersecurity is discussed has not been matched with new measurement techniques. In some organizations, the risk-based approach has gone a step further, with the application of the tools of enterprise risk management: once identified, cyber risk should be accepted, mitigated, avoided, or transferred. In most cases, though, the framework has left users to figure out a more developed measurement method on their own.

Filling in the Framework's gaps

Up to this point, the framework's impact on the practice of cybersecurity has been universal but diffuse. While it took a step toward discussion about cyber as risk-based, it did not provide the specific tools and use cases so that users—primarily cybersecurity professionals—could measure and manage that risk. When NIST next revises the framework, it should provide a means for users to measure cyber risk, including quantitative methods, which would allow organizations to assess the cost-effectiveness of their cybersecurity programs (a goal included at the outset by EO 13636). NIST is now working on better integrating the practice of cyber risk management into the broader disciplines of enterprise risk management with the goal of more closely aligning cybersecurity decision-making with overall organizational decision-making.[9] The internal reports it has begun producing start to address the challenge of making cyber risk a more quantitative discipline. A meaningful measurement of risk would mean organizations could put a value on reduced risk, just as they can measure cybersecurity spending. Comparing the cost of an activity with a reasonable measure of the benefit that will accrue from the activity is the way all organizations make decisions. Organizations using the tools currently provided by the framework cannot measure cost-effectiveness.

5 Cybersecurity Enhancements Act of 2014, 5 U.S.C. 272(c). The act adopted the language of EO 13676, which had resulted in the publication of the first version of the Framework in 2014 and mandated its continued development. Ibid., §272(e)(1)(A)(iii).

6 National Institute of Standards and Technology (NIST), Framework for Improving Critical Infrastructure Cybersecurity, version, 1.1, 2018. For information about the framework and how it functions, see the Cybersecurity Framework website.

7 In the first paragraph of its executive summary, the framework identifies cybersecurity risk as a category of risk, along with financial and reputational risks, that can impact an organization's bottom line. Ibid., p. v.

8 NIST Cybersecurity Framework, p. 4.

9 Kevin Stine, Stephen Quinn, Gregory Witte, and Robert Gardner, Integrating Cybersecurity and Enterprise Risk Management (ERM), NISTIR 8286, October 2020.

NIST on its own cannot create a cybersecurity cost-effectiveness measurement system; as discussed below, building such a measurement system will require addressing a variety of challenges and requirements. Nonetheless, a revised framework could build consensus around the importance of organizing a cybersecurity program around quantitative risk measurement, just as the existing one has supported a shift in talking about cybersecurity as a risk-based discipline. A revised framework would set forth the basic conceptual foundation for building models that would permit users to measure cost-effectiveness. NIST can build on the success of the framework by adding measurement and quantification as key instructions and core resources for its users.

Standardized Measurement for Cybersecurity Risk

Measurement requires a definable characteristic to measure and a standardized way to quantify that characteristic. For example, a measuring system for height first must identify height as a characteristic measured with units of length. If measurement is then done based on the length of one person's hand it is only useful only to people who know that person. A standardized unit of measure for length—like meters or feet—can allow comparison of height across a whole population.

A concept like cyber risk first must be defined before it can be standardized. Risk, defined in actuarial terms, relates to outcomes and their likelihood. The common way of describing cyber risk is by evaluating the consequences and likelihood of scenarios covering nearly every conceivable cyber incident. Methodologies for measuring the component parts of risk in quantitative terms combine analysis of the data on previous incidents and trends with estimates of impact or likelihood from experts, using techniques to reduce biases and subjectivity.

A new vision of the framework that can drive effective cyber risk measurement will require adopting a "yardstick strategy" in which measurement of cyber risk is standardized in a system that is used across many organizations.[10] In contrast with the current framework, such a revision would include direct guidance on techniques for collecting risk information in quantitative terms. The industry uses a variety of techniques but NIST will—correctly—prefer not to identify a specific methodology. More practical guidance on what quantitative cyber risk information looks like, how to collect it, and what operations it can legitimately be used for would give practitioners a basic roadmap for how to adopt a system that will fit into a larger standardized ecosystem.

These revisions should be consistent with the role NIST plays in the ecosystem. A revised framework need not mandate a single approach or replace the role of the private sector and independent standards-making bodies in creating the standards and conformity assessment mechanisms that would underpin a larger cyber risk measurement ecosystem.

Why Standardized Risk Measurement Matters

Like the example of having a system for measuring height that works across a population, a quantified measure of cyber risk would permit useful comparisons between organizations with otherwise very different characteristics and a way for them to compare themselves with a meaningful benchmark for risk treatment. The systems available to organizations should support robust and repeatable results based on measurements tested and verified by neutral arbiters like standards development organizations and national standards bodies. This would satisfy the desire of companies to know how they compare to their peers. The point is not to institute a public register of cyber risk—organizations may wish to avoid sharing details of their efforts to manage cybersecurity risk for good reasons related to liability and reputation—but to offer that greater certainty to organizations that they are measuring risk in a meaningful way.

10 The framework is used here as a stand-in for a larger effort; what is described will not come in one revision. See Kevin Stine, Stephen Quinn, Gregory Witte, and Robert Gardner, Integrating Cybersecurity and Enterprise Risk Management (ERM).

A cyber risk approach that had implemented such a "yardstick strategy" would also benefit national security. Policymakers currently lack a consistent way of measuring cyber risk at critical infrastructure companies. Without a uniform way to measure such risk, there is no easy way to assess adequacy or make comparisons regarding the preparedness of companies, groups of companies, or sectors. Developing a yardstick for measuring cyber risk would be the first step needed to facilitate direct comparisons, offer the opportunity to conduct meaningful preparedness exercises—a cybersecurity stress test for critical infrastructure—and ultimately allow for the creation of programs to incentivize risk management decisions that would support national security. While there are differences between discussing national security risk in cybersecurity and merely combining the risk assessed in a group of entities, the two are related.

Government-led Steps toward Standardized Risk Management

Building a conceptual model for cyber risk requires data to validate that the model reasonably reflects reality. The basic actuarial structure of the risk model for cybersecurity requires data collection from real-world examples about the range of impacts from cyber incidents, the characteristics of the victim organizations, and information about the nature of the incidents including technical information about the attacks. With sufficient data of this type and validated sampling techniques, analysts can generate more precise analyses of incidents in general; by type; and according to the industry, size, region, or other characteristic of a victim organization.

Today, this information is collected in haphazard ways and in irregular data stores. For example, all 50 states and the District of Columbia now have a data breach notification law.[11] This would appear to be a good source of data about cyber intrusions. In fact, states have adopted varying standards for when a breach triggers a notification requirement and failed to adopt a standardized set of data about breaches that a victim organization must report. In most cases states impose no requirement that the data be provided in a machine-readable format. As a result, the insights available in mandated breach reports at the state level are atomized and difficult—if not impossible—to analyze.[12]

U.S. government agencies and private companies have already done some of the work to define the type of information needed to build a better actuarial model and some commercial cybersecurity companies that have focused on using available data to assess cyber risk. These efforts provide a roadmap for an updated framework and a measurement system for cyber risk that will improve the landscape for decision-making in cybersecurity.

One of these key efforts to measure cyber risk emerged from a partnership between the insurance industry, the federal government, and academic researchers. In 2015, the U.S. Department of Homeland Security convened the Cyber Incident Data Analysis Working Group (CIDAWG), inviting cyber insurers, chief information security officers, and other cybersecurity professionals from various critical infrastructure sectors to participate. The group released three white papers that described what information firms should collect on cyber incidents,[13] the value

11 Data breaches are a subset of cyber incidents involving the loss of personally identifying information (PII) that can result in damage to individuals. If it affects a sufficient number of records under the differing state standards, it requires disclosure to these individuals. Cyber incidents may include other types of cyber events that result in losses. One example of an incident that typically has not triggered data breach notification is a successful ransomware attack, which traditionally encrypts a company's system without exfiltrating individuals' PII. For more on this distinction, see Verizon, 2020 Data Breach Investigations Report.

12 Privacy Rights Clearinghouse, Chronology of Data Breaches, undated.

13 DHS, Enhancing Resilience through Cyber Incident Data Sharing and Analysis: Establishing Community-Relevant Data Categories in Support of a Cyber Incident Data Repository, September 2015.

of collecting such data,[14] and how such data would be kept secure.[15] The white paper describing what information should be collected listed 16 categories of information about cyber incidents that "could be used to perform trend and other analyses by enterprise risk owners and insurers."

However, since the end of 2015 when the CIDAWG completed its work, the cyber insurance industry has not actively tried to create the contemplated data store. This is because its potential benefits have not outweighed the costs of setting it up. Based on the current state of underwriting, it seems no insurance company has collected all the CIDAWG-mandated data from incidents suffered in its portfolio, and most other organizations are unwilling to share information about adverse incidents voluntarily because concerns about liability continue to outweigh the potential benefits. In other words, the insurance companies are satisfied with the status quo.

Voluntary cyber information sharing is often held out as a solution to this problem, but it has demonstrably failed to provide benefits of the type needed to improve risk modeling. Part of the reason is that sharing incident information has costs: the cost of collecting and maintaining such information as well as the perceived cost associated with the possibility of the data's unintentional release. By offering a liability shield to stop potential claims based on privacy or other injury caused by sharing information—as the Cybersecurity Information Sharing Act of 2015 does—the liability concern appears to be addressed. There is no need to look further than the current state of voluntary cyber information sharing following the enactment of that law to see that organizations remain concerned about the unplanned release of such information to the market and would rather avoid the cost and risk associated with collecting and sharing incident information.[16]

Another concern is that the information needed to build these models is not intended for immediate incident response or other operational needs. The benefits from this sharing arise only when the entity working with the data can aggregate a lot of data over time to generate useful models for projecting losses. Accordingly, individual businesses will avoid sharing initially for this public-good purpose without a requirement to do so.

The CIDAWG, in addition to describing the data criteria needed for underwriting cyber incidents, highlighted the need for a data repository that could collect information about incidents while ensuring anonymity and confidentiality.[17] Such a repository would have to be extensive and cover organizations of a variety of sizes and in a variety of industries before robust analysis would yield meaningful results. Without mandates, the inherent benefits to an organization choosing to participate in such a data repository would be diffuse and unlikely to return the value of participation in the near term.

Private Sector-led Steps toward Standardized Risk Management

In addition to government-led initiatives, the private sector has also taken steps toward the quantification of cyber risk. For example, insurance companies offer cyber insurance to firms seeking to transfer their cyber risks. While the insurance industry often excels at quantifying and financializing risk, the increasing availability of cyber insurance has not yet translated into better measurement systems for cyber risk. Until very recently, cyber insurers have performed very simple underwriting before issuing cyber policies. That has been a result of the competitive economics of the business and exacerbated by the fact that the

14 DHS, Enhancing Resilience through Cyber Incident Data Sharing and Analysis: The Value Proposition for a Cyber Incident Data Repository, June 2015.

15 DHS, Enhancing Resilience through Cyber Incident Data Sharing and Analysis: Overcoming Perceived Obstacles to Sharing into a Cyber Incident Data Repository, December 2015.

16 Cybersecurity Information Sharing Act of 2015, 6 U.S.C.A. §1501 et seq. December 18, 2015.

17 DHS, Enhancing Resilience Through Cyber Incident Data Sharing and Analysis.

industry has found that the metrics currently used by their customers and potential customers do not correlate with cybersecurity risk. Those metrics—usually related to cybersecurity control implementation rather than explicit cyber risk metrics—provide only lagging indicators of cybersecurity and have not been found to correlate to the likelihood of an incident or an organization's resiliency should one occur. With the advent of the coronavirus pandemic, increased incidence of ransomware, and the major SolarWinds and Microsoft Exchange incidents, the increased claims and loss ratios due to the skyrocketing cyber threat in 2021 have created a renewed interest in better metrics for cyber risk and cyber defensive posture.

Some cybersecurity companies and consultants offer products and services that use publicly available data to assess cyber risk for organizations on a customized basis. The Open Group's Factor Analysis of Information Risk (Open FAIR) provides a technique to estimate cyber risk in financial terms by collecting information about the magnitude of potential losses and modeling likelihood as a balance between cybersecurity threats, vulnerabilities, and attacker behavior and the capacity of the organization's cybersecurity program. The modeling done by these consultants relies on a combination of information specific to the network being assessed and data sets about incidents across the economy. Some of these data sets are maintained by volunteers such as the VERIS Community Database (the database used by the team that assembles the Verizon Data Breach Investigations Report each year).[18] Commercial data sets of this type include one maintained by Advisen, which follows a similar approach of digesting public reporting of cyber incidents into their database.[19]

The industry leaders in measuring cyber risk currently use these data sources to provide analysis to clients. Publicly available data is inadequate but can reduce uncertainty to a limited degree. Those gains are difficult to replicate across an industry or sector, however, because of the lack of uniform metrics and validated data sets. This part of the cybersecurity services industry demonstrates that relying exclusively on market forces to spur the quantification of cyber risk is insufficient.

Companies attempting to measure cybersecurity risk in quantitative terms have struggled to succeed in the cybersecurity services market, but they have provided a roadmap for how to model risk using the limited available data. The growth in the use of the Open FAIR standard for measuring cyber risk,[20] research published last year by the Cyentia Institute providing insight into the ways companies of different sizes and in different industries experience cyber incidents,[21] and the work of independent analysts such as Doug Hubbard[22] have offered a proof of concept for a more robust and quantitative measurement of cybersecurity risk. The work of these experts has demonstrated that with better cyber incident data available, a yardstick measurement system for cyber risk is possible and beneficial. FAIR, for example, has seen a dramatic uptake over the last several years even as the practitioners using it remain a tiny fraction of the threat intelligence, vulnerability management, and control implementation segments of the cybersecurity market.

What about Qualitative Risk Management Techniques?

Some organizations that have moved in the direction of measuring their cybersecurity risk based on qualitative estimates of impact and likelihood will assert that they are already using risk to prioritize cyberse-

18 Available at GitHub.

19 The data set is available at Advisen, Cyber Loss Data.

20 Jim Hietala, What is Open FAIR™?, The Open Group, January 24, 2017.

21 Cyentia Institute, "IRIS 20/20 Report, May 2020, and "IRIS 20/20 Extreme Report, November 2020. Both use Advisen cyber incident data to draw conclusions about the scale of incidents in different industries, focused on the Fortune 500. The Extreme report looks at the impact on companies of different sizes in different industries handle the tail risk incidents, which are the worst among all those reported.

22 Douglas W. Hubbard and Richard Seiersen, How to Measure Anything in Cybersecurity, John Wiley, 2016.

curity decisions. Using a scale of low/medium/high or other similar ordinal terms, they may have used those estimates to create a risk matrix or heat map. By color-coding the different quadrants of the heat map, different risk scenarios are frequently classified on a stoplight basis: green, yellow, and red. Heat maps and stoplight charts are a standard tool for communicating risk based on ordinal measurement of its components. The result of such an exercise is a qualitative risk management assessment that can provide useful limited guidance on what risks should take priority.

Qualitative techniques can be valid and useful for executive reporting and providing guidance on prioritization but one of the mistakes made when using a qualitative system to measure cyber risk is investing time and resources in making ever more detailed heat maps based on qualitative estimates of impact and likelihood. Qualitative heat maps of risk do not support conducting additional analysis to reach more sophisticated conclusions. Because the original measurements are ordinal—that is, it is not possible to say that, for example, two "lows" equals a "medium" or two "mediums" equals a "high—doing math with them will not generate useful results. Fundamentally, a qualitative heat map suffers from that deficiency because it is an attempt to multiply such ordinal impact and likelihood measures. In addition, results untethered from terms that reflect real-world consequences (what does "medium" impact or likelihood mean?) make it impossible to make a spending decision informed by the range of impact reduction the model suggests. Put another way, how much is fixing a "yellow" worth?

Generating more valuable conclusions requires a quantitative approach and a quantitative approach requires a lot more data.

Recommendations

The framework sought to provide a cost-effective approach to assessing a user's cybersecurity program. Achieving this goal is critical to improving the practice of cybersecurity and safeguarding U.S. industry and critical infrastructure, but it depends on having useful metrics to measure cyber risk in quantitative terms. The recommendations provided here would create a set of building blocks that would make such measurement possible and even straightforward. They will require changes in the way the federal government collects, analyzes, and shares data about cyber incidents and further cooperation between the government agencies and the entire cybersecurity industry. In the end, many of the solutions that will improve cyber risk measurement will come from private entities, but only the federal government is well-positioned to set the conditions for that to be possible.

The president should issue an executive order charging the NIST with revising the Framework for Improving Critical Infrastructure Cybersecurity in a way that focuses on quantifying risk and charging the secretary of commerce with developing initial quantitative risk models from available data sources.

After seven years and one revision, the framework is due for a formal update. At the same time, with its internal reports to describe how cyber risk management can be integrated into a broader enterprise risk management framework, NIST is addressing the use of quantification methodologies to assess cyber risk. This is a welcome development.

NIST already has the authority to revise and update the framework without additional authorities, but through an executive order the president can direct resources that will broaden the scope and the reach of the framework. Similar to Executive Order 13636 that originally tasked NIST with developing the framework, a new executive order should include a presidential mandate to expand cooperation on cyber risk between the Departments of Commerce and Homeland Security. An executive order would also instruct the secretary of commerce to take initial steps to build cyber risk models using existing and available data (see the recommendation below on an open innovation competition for cyber risk). The expertise at the Commerce Department—not only at NIST but also

within the Economics and Statistics Administration on data collection and analysis—provides a good focal point to start the work necessary to improve the existing tools for cyber risk quantification. The goals should be to take the initiative and demonstrate how cybersecurity risk can currently be quantified and to lay out the beginning of a roadmap for future work in this area.

Congress should create a Bureau of Cyber Statistics at the Commerce Department with the authority to collect incident information from the private sector.

Enabling cyber risk measurement will require institutional changes in how the government is structured to drive change in how private organizations behave. The Cyberspace Solarium Commission's report last year recommended policies that would better prepare the United States to protect itself from a significant cyberattack as well as to respond and recover with more resilience should one occur.[23] One of the primary Solarium Commission recommendations was the creation of a Bureau of Cyber Statistics (BCS) to collect and acquire data about cybersecurity events and turn that data into actionable insights, datasets, and advisories for the public. The BCS would operate in a way similar to other dedicated statistical agencies in the federal government. It would acquire data related to cybersecurity in the United States that would be protected from disclosure to the public by law. In turn, it would share aggregated information, analysis, and data sets with the public, created using well-documented statistical techniques to ensure users cannot de-anonymize the data.

That basic exchange—collection of identifiable data in exchange for the release of useful and anonymized information—is used by a variety of statistical agencies already and has proven value. In addition to collecting information and aggregating it into anonymized data sets, the BCS would work with other cybersecurity-focused bureaus, including NIST, on identifying the measurement questions that underlie successful model building. Such modeling expertise should, in turn, help the cybersecurity industry shift its focus from control implementation to risk management. Such a capability in the federal government is critically needed not just to collect and share data, but also to lend credibility to the data and to data-guided risk management. The weight of the federal government will help counter the traditional resistance against rigorous risk management.

The Solarium Commission recommended that the BCS collect information from existing sources, including by purchasing commercial data. Others have emphasized that its initial focus should be on the federal government's systems. The proposal to limit the BCS to only data sets available on the market or from other government agencies appears to stem from an urge to make sure it has some time to build its capacity before taking on a larger mission.

There are, however, three major reasons that the BCS should have a complete mandate to collect information about incidents from across the United States, including private entities and state, local, tribal, and territorial (SLTT) governments. First, the portion of the Internet-linked information network in the United States that the federal government owns and controls is a tiny fraction of the whole. On top of that, the risk profile of the federal government is quite different from that of the private sector and even from that of SLTT governments. From a risk perspective, the federal government has the clear role as the guarantor of stability and predictability for the U.S. network as a whole (a similar role it plays in the economic and national security spheres). As a result, metrics based on the data collected from the federal networks alone would be less effective for risk management for the private sector and SLTT governments.

A second problem with limiting data collection to federal networks is that many of the critical mission elements on which it must deliver are more difficult to measure than the outcomes for private companies. In the private sector, the impact of a negative event—

23 Cyberspace Solarium Commission, Solarium Commission Report, March 11, 2020.

including cyber incidents—is measured almost exclusively in financial terms. Even second-order impacts such as reputational damage can be translated into consequences. This financialization of risk makes it much easier to design a yardstick measurement system because of the common units of measurement: dollars and cents. Limiting data collection to federal networks would miss the opportunity to develop and improve metrics relevant to the private sector.

Finally, should Congress create the BCS, it may be years before it revisits that law. A BCS that is left with only a federal function would have trouble accomplishing the broad mission suggested here. Rather than providing true public goods in the form of anonymized data sets that are useful to the cybersecurity industry and researchers, it would end up creating metrics that would have some impact in the federal government and, likely, would slowly seep out into the government contractor community. That model will not drive the type of change that the current crisis in cybersecurity demands.

Measurement of the distinctive marks of cyber incidents, methodologies for turning that data into models for projecting the cost of future incidents, and metrics that can be used in executive and budgetary decision making are the connective tissue between the design of a revised framework and the Solarium Commission's recommendations. The commission's insistence on improving the federal government's capacity to support better cybersecurity throughout the economy can provide the raw material for solutions that organizations can use to build measurement capabilities, adopt quantitative methodologies, and develop cybersecurity risk metrics. The initial goal within these organizations will be to link cyber risk assessments to the key outstanding budgetary issue: how much does a cybersecurity program reduce risk?

Congress should create a National Cyber Safety Board to investigate cyber incidents.

Actuarial modeling requires data about past incidents, but raw data does not explain how incidents take place. To build a model that will provide cybersecurity professionals some measure of guidance for what activities will reduce risk most effectively, that data must be linked to an understanding of these correlations. A strong investigative agency with the ability to compel the collection of evidence and report on the root causes of large-scale cyber incidents would make building those models possible.

The National Transportation Safety Board (NTSB) performs this function in the transportation sector, investigating accidents to ascertain their root causes, which has had a dramatic impact on safety in transportation.[24] NTSB's specific role as a government investigator is carefully constrained to ensure its investigations, where it has broad authority to pursue the facts and compel the production of evidence, remain independent of industry, the parties involved in an accident, and the regulators responsible for setting and enforcing rules on industry participants.

A National Cyber Safety Board (NCSB) could function in a similar way, with broad powers to compel cooperation with an investigation and a mandate to report on cybersecurity incidents but without authority to issue regulations or to create liability based on the behavior of those being investigated. In an NCSB investigation, the analysis would focus on the root causes of the incident in question. Collecting the evidence would be the first part of understanding how the incident unfolded and where failures resulted in injury. Over time, multiple investigations can reveal the correlations between cybersecurity activities and the incidents that result.

The connection between the proposed BCS and NCSB is the understanding of how failures in cybersecurity lead to incidents. The NCSB would not have the capacity to investigate all the incidents that occur each year. Given the broad potential societal impact of large-scale cyber incidents in terms of, say, the possible

24 For a thorough discussion of the pros and cons of such a proposal, see Scott Shackelford and Austin E. Brady, "Is it Time for a National Cybersecurity Safety Board?" Albany Law Journal of Science and Technology, January 12, 2018.

physical consequences of an attack on national critical functions or of a large number of Americans to have their personal information compromised in a large-scale data breach, the NCSB would focus on the largest incidents and only expand its focus to smaller incidents as resources allowed. Such a focus would provide model builders at the BCS and in the private sector with a better understanding of how the largest cyber incidents unfold and, if the scope of the investigatory authority mirrors that of the NTSB, the impact of the behavior of the people involved rather than a narrow focus on the information systems affected.

The federal government should run an open innovation grand challenge to demonstrate how quantitative models can lead to better cybersecurity outcomes.[25]

An open innovation grand challenge would have two components: the design of metrics to measure cyber risk and the development of a model that uses those metrics to accurately predict such risk. The results of the competition would provide insight into what data best enable predictive models. On that basis, the organizers would have a starting point for continued refinement of the most successful risk models. These metrics and models could be shared with the federal government to lay a foundation for the BCS.

The organizers should design and publicize the competition with a focus on enticing participants from industries such as insurance and cyber defense as well as academics and other risk professionals—potentially in cross-disciplinary teams. To encourage participation, all the teams would have the opportunity to commercialize their methods after the competition. The competition could also demonstrate the benefit for private-sector companies of sharing their incident data with a trusted third party like the BCS.

The competition could be run by any of the federal government agencies that have been given the authority to do so under the America Competes Act. Perhaps the most obvious candidate would be the General Services Administration's (GSA) Challenge.gov program. The GSA would likely benefit from support from cybersecurity agencies like the Cyber and Infrastructure Security Agency within the Department of Homeland Security, and the Department of Commerce, where existing statistical agencies would provide a useful template for the creation of the BCS at the Commerce Department as suggested by the Solarium Commission and as recommended above.[26] In the interest of building enthusiasm for the competition, the federal government might also seek to work with outside partners, including philanthropic donors to boost a potential prize pot and industry organizations to encourage private-sector participation.

Conclusion

The Framework for Improving Critical Infrastructure Cybersecurity has played an important role in changing the way cybersecurity is conceptualized, discussed, and understood—in the cybersecurity industry and among organizational leaders not well-versed in cybersecurity. Conversations about cybersecurity now revolve around risk and the management of it. At the same time, the framework in many ways is more abstract than practical, and it leaves users with too much latitude and not enough guidance on how to build toward a risk management approach to cybersecurity with increasing levels of sophistication over time. As a result, the question of what tools and models users could adopt to make their risk concrete was left unanswered. It is time to make changes in collecting, analyzing, and disseminating information to spur the innovation necessary to improve cybersecurity as a discipline.

25 This section draws on a recommendation already published as part of GMF Digital's #Tech 2021—Ideas for Digital Democracy. See Adam Bobrow, Launching a Cyber Risk Grand Challenge, November 2020.

26 Solarium Commission Report, p. 78.

This work represents solely the opinion of the author and any opinion expressed herein should not be taken to represent an official position of the institution to which the author is affiliated.

About the Author(s)

Adam Bobrow is the founder and CEO of Foresight Resilience Strategies, a strategic consultancy helping small and medium-sized organizations adopt a risk management approach to cybersecurity. Before starting Foresight, he served at the White House Office of Science and Technology Policy and at the Department of Commerce. Prior to joining the Obama administration, Adam had a career as a consultant and lawyer in the private and public sectors focused on the U.S.-Chinese relationship, China's legal development, international trade, and intellectual property. He received his JD from Washington University in St. Louis and his undergraduate degree from Georgetown University.

Acknowledgments

I received tremendous assistance from members of the cybersecurity community within and outside of government in refining the ideas in this paper. A cadre of readers offered essential critiques and suggested improvements in the manuscript. I am particularly indebted to assistance from officials at the Departments of Commerce and Homeland Security, staff of the Cyberspace Solarium Commission, congressional staff, and individuals in the insurance industry. Specific thanks are also due to the William and Flora Hewlett Foundation for sponsoring this work, GMF for its stewardship of this project over the past year, Auburn University's McCrary Institute, which was the original home of this project, the leadership of the Society of Information Risk Analysts, and Eli Sugarman and Ian Wallace.

About GMF

The German Marshall Fund of the United States (GMF) is a non-partisan policy organization committed to the idea that the United States and Europe are stronger together. GMF champions the principles of democracy, human rights, and international cooperation, which have served as the bedrock of peace and prosperity since the end of World War II, but are under increasing strain. GMF works on issues critical to transatlantic interests in the 21st century, including the future of democracy, security and defense, geopolitics and the rise of China, and technology and innovation. By drawing on and fostering a community of people with diverse life experiences and political perspectives, GMF pursues its mission by driving the policy debate through cutting-edge analysis and convening, fortifying civil society, and cultivating the next generation of leaders on both sides of the Atlantic. Founded in 1972 through a gift from Germany as a tribute to the Marshall Plan, GMF is headquartered in Washington, DC, with offices in Berlin, Brussels, Ankara, Belgrade, Bucharest, Paris, and Warsaw.

G | M | F

Ankara • Belgrade • Berlin • Brussels • Bucharest
Paris • Warsaw • Washington, DC

www.gmfus.org