An Effects-Centric Approach to Assessing Cybersecurity Risk

A CISSM Report | March 2019

Charles T. Harry & Nancy W. Gallagher

SCHOOL OF PUBLIC POLICY
CENTER FOR INTERNATIONAL & SECURITY STUDIES AT MARYLAND

Introduction

Faced with rapidly growing cyber threats, organizational leaders, and government officials cannot reliably secure all data and digital devices for which they are responsible. The best they can do is conduct strategic risk management. That requires a systematic way to categorize potential attacks and estimate consequences in order to set priorities, allocate resources, and mitigate losses.

The 2018 U.S. National Cyber Strategy[1] holds government officials accountable for doing cyber risk management based on the National Institute of Standards and Technology's (NIST) Cybersecurity Framework and recommendations from not-for-profit organizations such as the Center for Internet Security (CIS) and ISACA. Yet, none of these policy documents and best practice guides actually provide the necessary analytical tools. As a result, public agencies, private companies, and non-profit groups that try to do risk assessment often feel overwhelmed rather than empowered to make strategic cybersecurity decisions.

The Center for International and Security Studies at Maryland (CISSM) has developed an analytical framework that provides four essential building blocks needed to satisfy the principles in the NIST Standard Framework and other best practice guides:

1. A standardized system for classifying cyber threats and events by their effects.
2. Tools to associate organizational functions with IT topologies.
3. Algorithms to assess the severity of disruptive and exploitative cyber events.
4. A method to understand the integrated nature of risk across different parts of a simple organization, major divisions of a complex organization, or interconnected organizations in a complex system.

These building blocks can be combined in different ways to answer critical questions, such as:

- What is the range of cyber risks to different types of organizations?
- Which threats pose the greatest risk to a specific department or organization?
- How could an attack on one part of an IT network affect other organizational functions?
- What is the accumulated risk across a critical infrastructure sector or geography?

Using a comprehensive, consistent, and repeatable method to categorize and measure risk can enhance communication and decision-making among executives who make strategic decisions for organizations and their IT staff with day-to-day responsibility for cybersecurity. It can facilitate cooperation between public officials and private industry who share responsibility for different components of national critical infrastructure. It can inform media coverage and public debate about important policy questions, such as which decisions about cybersecurity should be purely private decisions, whether government should incentivize or mandate certain cybersecurity choices, and when a cyber attack warrants some type of military response.

[1] https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Cyber-Strategy.pdf

Understanding Cyber Threats

Media reports around the world warn constantly about steadily growing cyber threats. The headline of a *Newsweek* article read, "U.S. Hit by 77,000 Cyber Attacks in 2015 – a 10 Percent Jump."[2] The IT Governance compilation of data breaches reported 3.1 billion records leaked in 2016.[3] And, the Online Trust Alliance's annual report found that cyber attacks doubled in 2017.[4]

The United States intelligence community now puts cyber attacks at the top of its threat assessment list, above weapons of mass destruction, proliferation, and terrorism. The 2017 U.S. National Security Strategy directs government officials to identify and prioritize cyber risks across six key areas – national security, energy and power, banking and finance, health and safety, communications, and transportation. The 2018 National Cyber Strategy added a seventh key area: information technology itself.

These documents also underscore the U.S. government's willingness to use the full range of military options at its disposal to prevent, defeat, or retaliate against serious cyber threats or attacks without specifying what kind of potential or actual cyber events would actually warrant such action. The National Security Strategy threatens to impose "swift and costly consequences on foreign governments, criminals, and other actors who undertake significant malicious cyber activities."[5] For the first time, the National Cyber Strategy explicitly authorizes the use of offensive cyber military options "to prevent, respond to, and deter malicious cyber activity." The 2017 Nuclear Posture Review even threatens to use nuclear weapons if there is a "significant non-nuclear strategic attack" on critical infrastructure in the United States or allied countries.[6]

This raises important questions about how to classify cyber threats, assess risks, and determine whether a cyber event that affects some type of critical infrastructure is "significant" enough to warrant some type of military response. Unfortunately, the same words are often used for undesirable cyber events that range from inconsequential to catastrophic. For example in 2016, the Japan Times ran a story claiming that Japanese networks were hit with over 128.1 billion cyber attacks in 2016 alone, when much of that activity had no negative effect.[7]

Confusion about the Sony Hack of 2014 led Senator John McCain to call it an act of war:

> "[t]he president does not understand that this is a manifestation of a new form of warfare....When you destroy economies, when you are able to impose censorship on the

[2] http://www.newsweek.com/government-cyber-attacks-increase-2015-439206.

[3] https://www.itgovernance.co.uk/blog/list-of-data-breaches-and-cyber-attacks-in-2016-1-6-billion-records-leaked/

[4] https://otalliance.org/system/files/files/initiative/documents/ota_cyber_incident_trends_report_jan2018.pdf

[5] https://www.whitehouse.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905.pdf

[6] https://media.defense.gov/2018/Feb/02/2001872886/-1/-1/1/2018-NUCLEAR-POSTURE-REVIEW-FINAL-REPORT.PDF.

[7] https://www.japantimes.co.jp/news/2017/02/08/national/crime-legal/cyberattacks-targeting-japan-networks-hit-record-128-1-billion-2016/#.W97o0pNKhD8

> world and especially the United States of America, it's more than vandalism. It's a new form of warfare that we're involved in, and we need to react and react vigorously."[8]

Such pronouncements are unhelpful to officials who must categorize motivations and seek policy prescriptions that appropriately deal with the situation. In a 2015 House hearing on global cyber threats, former National Security Agency Director Admiral Rodgers warned, "Terminology and lexicon is very important in this space… I'll hear people throw out attack, act of war, [when] that's not necessarily … how I would characterize the activity that I see."[9]

Three years later, there still is no standard way of differentiating among various types of threats. Instead, a confusing array of cyber classification systems are used for different purposes. Some classification systems are based on phases of the hacking process or types of hackers (criminal, hacktivist, nation-state, etc.). Others categorize by attack vectors, vulnerabilities, techniques (spearphish, malware, etc.) or technology targeted (SCADA systems, cloud servers, etc.).

Since risk is a function of probability and consequences, the most relevant distinction involves the effects of a cyber event, not who did the deed or what tools and techniques they used. Unfortunately, the original effects-based approach developed by Howard and Longstaff (1998) and the more recent one by Kjaerland (2005) do not meet basic criteria for classification systems, including that every item to be classified fits into one, and only one, category.[10]

Each categorization scheme can illuminate a certain aspect of cybersecurity at a particular point in time. Yet none offers a comprehensive classification system that can remain useful as threat actors, technology, hacking techniques, and other aspects of the problem change over time.[11] This impedes communication, inhibits comparison, and complicates cumulative research. The multiplicity of classification systems also presents a major problem for organizational leaders and policymakers who must make strategic decisions about risk management across a wide variety of interconnected IT devices and networks that could be directly or indirectly affected by many different types of threat actors using attack vectors, vulnerabilities, and techniques that evolve rapidly over time.

Current Approaches to Risk Estimation: Useful, but Insufficient

Given the frequency and diversity of cyber events, trying to prevent all cyber attacks would be expensive and impossible. Therefore, organizational leaders and policymakers need some way to prioritize what high-consequence events they need to prevent at all costs, what mid-range ones

[8] https://www.huffingtonpost.com/2014/12/21/sony-north-korea-war_n_6362454.html
[9] House of Representatives Session titled "Cyber Security Threats," September 2015, https://www.c-span.org/video/?328021-1/hearingworldwide-cybersecurity-threats
[10] Howard, J. and Longstaff, T. (1998) "A Common Language for Computer Security Incidents," Technical Report, Sandia National Laboratories, and Kjaerland, M., (2005) "A taxonomy and comparison of computer security incidents from the commercial and government sectors". Computers and Security, Vol 25 pp 522–538.
[11] See Harry, C and Gallagher, N, "Classifying Cyber Events: A Proposed Taxonomy," *CISSM Working Paper* (February 2018) for a fuller assessment of other taxonomies.

they should try to mitigate, and what low-consequence ones they can live with easily. A risk-centric approach lets them be strategic about mitigation options and resource allocations.

The first major effort to provide guidance on cybersecurity risk assessment was part of the Information Security Management System published by the International Standards Organization and the International Electrotechnical Commission in 2005. ISO/IEC 27001 (updated in 2013) directs managers to systematically assess their organization's information security risks by evaluating threats, vulnerabilities, and impacts, but it does not specify how to do that. The same is true for the OECD's *Digital Security Risk Management for Economic and Social Prosperity*.[12]

A range of private sector efforts, including ISACA's COBIT 5 standard and the Center for Internet Security's security controls, have also framed risk assessment as central to effective cyber defense. The number of industry standards, guides, and recommendations for risk assessment grew substantially in the past decade in response to the rising volume, veracity, and magnitude of cyber incidents. The range of guidance, which was often disperse or subject to its own nomenclature, was eventually grouped and referenced under Obama administration efforts to specify cybersecurity best practices for critical infrastructure.

The NIST Cybersecurity Framework, first released in 2014 and revised in April 2018, is an umbrella document that serves as the standard for cybersecurity in the U.S. Executive Branch. It is also widely used by other organizations. It covers five categories of functional activities to improve cybersecurity. Risk identification and management are in the Identify (ID) category.

The NIST Standard Framework provides principles for systematic risk assessment including:

- Assess the potential consequences that could result from specific threats.
- Assess the realistic likelihood of those threats to engineer the consequences identified.
- Assess risk by asset, organization, function, and integrated risk between firms.
- Generate scenario-based use cases in estimating risk.
- Define a consistent approach to be used across the organization.

The Trump administration's 2018 U.S. Cybersecurity Strategy puts risk-based decision making at the center of its efforts to secure federal networks and to reduce threats to critical infrastructure. It aspires to "develop a comprehensive understanding of national risk by identifying national critical functions and will mature our cybersecurity offerings and engagements to better manage those national risks."[13] It also intends to "hold department and agency heads accountable for managing the cybersecurity risks to systems they control, while empowering them to provide adequate security."[14]

[12] https://www.oecd-ilibrary.org/science-and-technology/digital-security-risk-management-for-economic-and-social-prosperity_9789264245471-en

[13] US Cybersecurity Strategy, September 2018 https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Cyber-Strategy.pdf

[14] https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Cyber-Strategy.pdf

The NIST Standard Framework does not provide the tools needed to assess risk in a way that satisfies its principles. Instead, it directs users to the aforementioned best-practice guides and several NIST publications (e.g. SP 800-30 and SP 800-53) for help assessing risk in a way that fits their organization's mission and IT infrastructure. A one-size-fits all approach to cyber risk assessment would be ineffective. Yet, this assortment of best practice guides also has at least three major shortcomings. Thus, government officials and organizational leaders are being held accountable for doing something that they have not actually been empowered to do very well.

One problem involves oversimplification. These best practice guides direct IT managers to estimate potential impacts as if the consequences of an attack on a particular type of device (desktop computer, router, SCADA system) would be the same regardless of what type of event occurred. For example, NIST SP 800-30, "Risk Management Guide for Information Technology Systems," suggests estimating the severity of an actual or potential incident as having a High/Medium/Low impact on the confidentiality (C), integrity (I), and accessibility (A) of the targeted device. Risk assessment involves entering judgments about impact into a grid with columns for the C, I, and A dimensions and rows for different types of devices, such as workstation computers, routers, and servers. NIST SP 800-30 provides no standards for qualitative judgments, nor a quantitative method to calculate scores.

A grid based on the CIA triad is a convenient way to visualize and assess different types of cyber risks across a complicated IT system, but different types of cyber attacks on the same device could have very different consequences. For example, stealing information from a file server could have "High" impact on confidentiality and Low impact on integrity and accessibility, while interfering with that file server's operation might have Low impact on confidentiality and integrity, and a Medium or High impact on accessibility. Moreover, a ransomware attack might make the server unavailable for weeks, while a denial of service attack might make the data inaccessible for a few minutes to hours. Both attacks affect the accessibility of the system but have differing durations, thereby engineering different effects. They would be scored the same in a CIA assessment of effect, even though the range of impact is quite substantial.

A second shortcoming involves disaggregation of complex systems. NIST SP 800-30 and similar guides assess how specific threats to individual IT assets could directly harm an organization's operations and data security. They do not conceptualize individual devices as part of an interconnected IT system where a disruption in one component could indirectly impede operations supported by other parts of the IT system. Yet, the same type of attack on the same type of device could have different effects depending on where the affected device fit into the organization's overall IT system. Information stolen from an individual workstation could be embarrassing for the employee targeted, but it would probably not have a high C, I, or A impact for the whole organization unless that workstation belonged to a particularly important person or somebody who worked with very sensitive corporate information.

Assigning the same low, medium, or high-risk scores for all devices of the same type irrespective of what role each plays in the organization's IT system produces a laundry list of vulnerabilities without key information needed to set priorities. It might be, for example, that a problem with router A should be prioritized over a different vulnerability in router B because router A is a central node supporting many devices used to manufacture the organization's most important product, while router B supports a subset of computers used for human resource functions.

The third weakness involves insufficient differentiation of effects. The recommended guides typically treat adverse effects of cyber events as a diverse set of bad consequences without differentiating among different types of effects about which different stakeholders might be most concerned. For example, ISO 27001 lists various things that organizational leaders want to avoid, including revenue lost, reputational damage, decline in stock value, or higher insurance costs. Each of these potential harms, and more, are relevant for risk assessment. Yet, some should matter more to an organization's Chief Information Officer than to its Chief Executive Officer. Different stakeholders outside of an organization will also want some way to differentiate among various types of adverse effects, whether they are selling cyber insurance, relying on the organization for key goods or services, or overseeing the critical infrastructure sector to which this organization belongs. NIST SP 800-53 directs users to consider "tiers of impact." It does not provide a way to assess how vulnerabilities in one part of an interconnected IT network can lead to various types of bad consequences in other parts of the organization's IT system, its supply chain, or its community.

Some cybersecurity companies and not-for-profit organizations offer proprietary techniques for customized risk management, but publicly available information suggests that they are also insufficient. For example, the Factor Analysis of Information Risk (FAIR) Institute promises that its approach can "enable an organization to cost-effectively achieve and maintain an acceptable level of loss exposure." Its method helps business leaders and risk professionals explicitly define what asset, threat, and loss mean for their organization. It also directs them to identify metrics for asset vulnerability, threat development, and actual or potential financial losses incurred if threat actors take advantage of vulnerable assets in some way.

While this approach applies a more thorough and repeatable method for an individual organization to quantify risk, its allowance for non-standardized effects classification, its focus on financial loss vice other effects, and its non-systems-based measurement of impact create significant shortcomings in its application to assess interconnected and complex impacts stemming from a range of cyber-attacks. For example, while a firm might value a specific database as an important element of its packing and delivery service, the database's relationship with other devices used in the conduct of that service (e.g. desktop computer and routers) are also important to include when measuring risk to that operational function. The FAIR standard does not provide a way to do this because its scoring functions do not account for the number and interrelationship between those devices.

The FAIR approach also does not reflect how cyber events can have a range of primary, secondary, or second order effects on various things of great value to stakeholders outside the organization. From the perspective of a business leader using the FAIR approach, it would not be cost-effective to patch a vulnerability in software running on an administrative assistant's workstation if the financial loss associated with a cyber attack on that machine seemed acceptable. Such calculations could change if one took into account the possibility of an attacker moving laterally through the IT system from its point of entry to a server storing some of the company's trade secrets or a SCADA system operating a critical piece of equipment. That could entail much more significant financial losses and other types of bad consequences for the targeted organization, plus serious side effects for anyone who depended on the goods or services that organization provided.

In short, none of the common digital security risk assessment methodologies offers a straightforward way for organizational leaders and their top IT staff to do the type of holistic review that the NIST Standard Framework recommends.

The CISSM Framework: A Holistic Approach to Assessing Integrated Cybersecurity Risk

To address problems created by vague language, a confusing multiplicity of categorization schemes, and shortcomings with available risk assessment approaches, CISSM has developed a framework with four main components:

1. A standardized taxonomy for classifying cyber threats and events by their effects.
2. Tools to associate organizational functions with IT topologies.
3. Algorithms to assess the severity of disruptive and exploitative cyber events.
4. A method to understand the integrated nature of risk across different parts of a simple organization, major divisions of a complex organization, or interconnected organizations in a complex system.

Component 1: Effects-Based Classification System

CISSM's risk assessment framework starts with an effects-based taxonomy that divides all cyber events into two main categories—disruptive or exploitative—depending on whether the primary objective is to interfere with some function of the targeted organization or to steal information.

We define a cyber event as the result of any single unauthorized effort, or the culmination of many such technical actions, that *engineers, through use of computer technology and networks, a desired primary effect on a target.* This definition does not include many things that are often covered by the vaguer term "cyber incident," such as probes for vulnerabilities, failure to follow proper cybersecurity procedures, or use of social media accounts to spread false information about a political candidate.

If the same campaign produces several different types of effects, we treat it as multiple events that occurred simultaneously or sequentially. For example, during the 2014 Sony Pictures hack, the threat actor leveraged various tactics to access an application server from which it could compromise and exfiltrate e-mail, memos, and other organizational data. It also deleted data on corporate e-mail servers, thereby disrupting Sony's internal operations. The campaign, which likely involved thousands of person-hours to access and propagate through network, is classified by its end results: exploitation and disruption.

Disruptive Events

Disruptive effects can be classified into five sub-categories depending on what part of an organization's IT infrastructure is most seriously impacted.

Message Manipulation interferes with a victim's ability to accurately present or communicate its "message" to its customer base or other audience. These attacks include the hijacking of social media accounts, such as Facebook or Twitter, or defacing a company website by replacing the legitimate site with pages supporting a political cause.

External Denial of Service events are executed from devices outside of the target organization's network to degrade or deny the victim's ability to communicate with other systems. Many types of Distributed Denial of Service (DDoS) attacks would fit into this category, including ICMP flood, SYN flood, or ping of death. A Border Gateway Protocol (BGP) hijack that diverted Internet traffic away from a targeted organization's website would also fit in this category.

Internal Denial of Service events are executed from inside a victim's network to degrade or deny access to other parts of the IT network. For instance, an attacker who gained remote access could move laterally inside an organization's network to reset a core router to factory settings, preventing devices inside the network from communicating with each other. He could also install malware on a file server and disrupt data sent to and received from user workstations.

Data Attack events manipulate, destroy, or encrypt data in a victim's network. Common techniques include the use of wiper viruses and ransomware. Using stolen administrative credentials to manipulate data and violate its integrity, such as changing grades in a university registrar's database, would fall into this category, as well.

Physical Attack events use IT components, such as SCADA systems, to manipulate, degrade, or destroys physical systems. Current techniques used to achieve this type of effect include the manipulation of Programmable Logic Controllers (PLC) to open or close electrical breakers, leading to a de-energizing of that portion of the grid, or the utilization of user passwords to change settings in a human machine interface so that a blast furnace overheats and is destroyed.

Exploitive Events

Exploitative events occur when the hacker's primary motivation is to steal customer data, intellectual property, classified national security information, or sensitive details about the organization itself. The CISSM taxonomy classifies exploitative events by the part of an entity's IT infrastructure from which the malicious actor steals the information.

Exploitation of Sensors events occur when data is stolen from a peripheral device, such as a credit card reader, smart TV, or baby monitor. For example, in 2013, the Target corporation had thousands of their Point of Sale (PoS) devices compromised, leading to the loss of over 40 million customer credit card numbers.[15]

Exploitation of End Host events steal data stored on user's desktop computers, laptops, or mobile devices. Common tactics currently used include sending a malicious link for a user to click or leveraging compromised user credentials to log in to an account.

Exploitation of Network Infrastructure events involve the compromise of data through direct access to networking equipment such as routers, switches, and modems. In one 2018 example, over 500,000 routers worldwide were infected with VPNFilter malware which maintained access to devices through the compromise of user credentials and left open the potential for information to be hijacked.[16] The access to network devices allowed users to siphon information about where devices were located, potentially privileged areas of the network segmented off by the administrator, or other technical data useful in selling on the dark web or used by the hackers to expand their access.

Exploitation of Application Server events occur when malicious actors use a misconfiguration or vulnerability to gain access to data in a server-side application (e.g. a database) or on the server itself. The hacker in the 2015 Office of Personnel Management data breach used a SQL injection to access millions of records with sensitive information about current and former government employees. This category also includes the theft of data from Sony Pictures achieved when the hacker gained direct access to an e-mail server with organizational correspondence.

Exploitation of Data in Transit events occur when hackers acquire data moving between devices. For example, unencrypted data might be acquired as it is sent from a PoS device like a credit card reader to a database, or when somebody makes a purchase over the Internet from their laptop through an unsecured wireless hotspot at a local coffee shop.

Differentiating among various types of cyber events using this taxonomy enables cyber risks to be discussed, measured, and addressed in a more precise manner. For example, a substantial percentage of the 77,000 incidents in 2015 that the *Newsweek* headline termed "cyber attacks" were actually incidents without effects. The federal reporting requirement which produced this statistic requires notification when acceptable use policies are violated by government employees or standard security practices are ignored.[17] The IT Governance blog's 3.1 billion statistic refers to its own running count of individual records stolen through cyber espionage, not how many or what type of exploitative events occurred. The Online Trust Alliance's assessment that cyber incidents had doubled from 2016 to 2017 included unauthorized access without effects, plus various exploitative events, disruptive events, and other "activities causing financial or reputational harm."

[15] Shue, et al., "Breaking the Target: An Analysis of Target Breach and Lessons Learned," https://arxiv.org/pdf/1701.04940.pdf

[16] "VPNFilter: New Router Malware with Destructive Capabilities," May 2018, Symantec, https://www.symantec.com/blogs/threat-intelligence/vpnfilter-iot-malware,

[17] http://www.us-cert.gov/government-users/reporting-requirements.

Such data collection efforts can be valuable, but they can also be misleading if the definitions and data collection methods are not understood. Important distinctions among types of cyber events are blurred, trivial incidents can count as much as events with serious consequences, and some types of cyber events do not receive attention because they fall outside of whatever categorization system is being used.

There is no public dataset of cyber events that includes the full range of exploitative and disruptive events covered by the CISSM taxonomy. To create one, CISSM researchers used systematic web searches of English-language news sources from January 2014 to December 2016 to identify 2,431 cyber events whose effects were described in enough detail that they could be categorized. This dataset is far from complete: many cyber events are too trivial or too sensitive to get media coverage. Nevertheless, an analysis by sector produced some insights that could improve risk assessment. For example, the Government and Professional Services sectors suffered nearly half of all attacks recorded during this period, while a few sectors, including agriculture and construction, were largely immune. The type of cyber event also varied by sector. They almost exclusively involved data theft in the hotel, food service, and retail sectors, while they were almost evenly divided between the exploitative and distributive categories in the Government, Information, and Arts and Entertainment sectors.[18] Moreover, exploitive attacks in the retail sector focused on acquisition of data in PoS devices (e.g. Exploitation of Sensor), whereas hospitals saw most of their attacks aimed at patient databases (Exploitation of Application Server) or aimed at denying access to those devices through ransomware or wiper attacks (Data Attack).

Component 2: Tools to Associate Organizational Functions with IT Topologies

The consequences of a cyber event will differ depending on how the targeted device connects to other IT assets and how these networks support specific functions that an organization must fulfill to carry out its mission. For example, if an organization has two employees enter payroll information on their laptops, which connect to an application server through a router, somebody filling out a CIA grid might estimate the impact of an attack on the application server as high and the impact of an attack on the employee laptops as low. Yet the laptops, router, and the application server work together to support the payroll management function. Therefore, risk assessment should consider multiple different attack scenarios that could disrupt payroll processing, or other important organizational functions, rather than just thinking in terms of potential impacts on individual devices.

CISSM's approach starts by identifying user-defined organizational functions, their underlying IT components, and network connections across, into, and out from the organization. If the person using CISSM's analytical framework has detailed information about a specific organization's functions and their supporting IT infrastructures, that can be used to create a customize topology. If not, general knowledge about organizations of that type can be used to create a more generic map representing how they use information technology to carry out different activities required to produce goods or provide services. For example, most

[18] See Harry and Gallagher, "Classifying Cyber Events: A Proposed Taxonomy."

organizations use some common types of IT systems to communicate externally and internally and to store personnel or customer records. They typically will also have some IT systems that support functions specific to their organizational mission. A university will rely heavily on IT for research and educational purposes, while a manufacturing plant will use it more to operate machines and maintain its supply chains. The CISSM framework divides an organization's network topology into node-edge graph models tied to functions, called "Cyber Strands."

The next step in the risk assessment process asks an organization's IT managers, or whoever else is performing the risk assessment, to use scenarios to make judgments about how susceptible the various devices that comprise different Cyber Strands are to the ten types of cyber events in the CISSM taxonomy. The likelihood of an event is typically treated as a function of the number or type of Common Vulnerabilities and Exposure (CVE) records associated with the hardware or software tied to the device being analyzed. However, the likelihood of a specific type of cyber event happening to a specific part of an organization's IT network is also a function of the humans who have legitimate reasons to interact with that software or hardware, as well as the motivations and capabilities of hackers who might try to misuse it.

Scenario-based analysis helps IT managers and other risk analysts get beyond the most obvious threats and vulnerabilities by thinking like a clever hacker considering different ways to achieve some desired effect. That hacker might use a known vulnerability in the same way it had been used before, leverage it for a different purpose, or exploit other unidentified vulnerabilities, including insider access. The scenarios also encourage risk analysts to consider not only what is possible, but also what is plausible. Yet, there are many different kinds of hackers and many different types of vulnerabilities, but only a highly motivated, extremely skilled, well-resourced threat actor is likely to be willing and able to carry out certain types of attacks.

This step in the analysis produces a manageable list of plausible cyber events of particular concern for a given organization. The next step in risk assessment is to estimate potential consequences so that organizational leaders and policymakers can prioritize prevention and response for high-consequence events even if they are relatively unlikely compared with less consequential ones.

Component 3: Assessing Severity of a Cyber Event's Primary Effect

Estimating impacts resulting from cyber attacks mostly remains a subjective determination, even when policy makers must determine whether an incident warrants a governmental response. Presidential Decision Directive 41, the Obama administration policy that first addressed this problem, assessed severity using a six-point index based on qualitative judgments about the likelihood and level of damage to American interests. This scale is a useful reminder that many cyber events are not worthy of a government response, and that only the most severe are likely to have a significant effect on "public health or safety, national security, economic security, foreign relations or civil liberties."[19] Because it provides no standard and repeatable way to make judgements about likelihood and level of effects across these very different types of interests,

[19] https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Cyber%2BIncident%2BSeverity%2BSchema.pdf.

though, it could foster intense disagreements about if and how the government should respond to a severe attack, impeding consequence management or fostering over-reaction. The PPD-41 Cyber Incident Severity Schema also lacks any methodology for estimating likely effects of specific potential attacks on specific kinds of organizations, making it irrelevant for risk assessment.[20]

To address these problems, the CISSM framework differentiates among primary, secondary, and second-order effects. Primary effects are the direct impacts to the target organization's data or IT-enabled operations. Cyber events can also cause secondary effects to the organization, such as the financial costs of replacing equipment damaged in an attack, a drop in the organization's stock price due to bad publicity from the attack, or a loss of confidence in the organization's ability to safeguard confidential data. And, they can cause second-order effects on individuals or organizations who rely on the targeted organization for some type of goods or services. These could include effects on the physical environment, the supply chain, or even distortions an attack might have on an individual's attitudes, preferences, or opinion deriving from the release of salacious information.

The CISSM framework makes an initial estimate of severity based on the primary effects of actual or hypothetical cyber events, then determines secondary financial impacts and second order effects based on the direct impacts to the network's ability to function or to the amount of data lost. For example, if 100,000 customer records are compromised in a data breach, the calculation of primary effect provides a first estimate of the relative scale, which then can be paired with estimated per unit costs to estimate the secondary effect to the institution.

The severity of a cyber event's primary effect is assessed by the Cyber Disruption Index (CDI) or the Cyber Exploitation Index (CEI). Users can estimate effects from an event that has actually occurred or a specific attack scenario they want to evaluate using algorithms that generate scores on a 0 to 1 scale. The specific algorithms are proprietary, but the factors of analysis for both are included below. Severity scores can be calculated more precisely, and with greater confidence, by algorithm users who have granular knowledge about the IT network and organization impacted, while order of magnitude severity can be estimated with more generic knowledge.

Estimating the Severity of Disruption: Cyber Disruption Index (CDI)

The actual or potential consequences of a disruptive event are a function of its scope, magnitude, and duration. For example, a ransomware attack against a single device in the front office of a multi-billion-dollar organization might not be as consequential as a denial of service attack against a core router that prevents all users from communicating on the network, especially if the company could quickly restore the data using a back-up copy but could not resume normal internal communications for many hours.

[20] Harry, C and Gallagher, N, "Categorizing and Assessing the Severity of Disruptive Cyber Events," *CISSM Policy Paper* (April 2017), at: http://www.cissm.umd.edu/publications/categorizing-and-assessing-severity-disruptive-cyber-incidents.

The Cyber Disruption Index utilizes a graph method to integrate information about the number and centrality of devices affected (scope), the negative impact on organizational function of each asset affected (magnitude), and the duration of the effect into a single calculation. The CDI user would make an explicit judgment about whether the number and importance of IT devices effected was insignificant (0.2), minimal (0.4), significant (0.6), massive (0.8), or total (1.0). They would use the same scale for magnitude by assessing impact on productivity. They would assess duration using units of time they considered insignificant, minimal, or worse for the organization in question—e.g. seconds, minutes, hours, days, or months of downtime for the affected part of the organization. Multiplying the scope, magnitude, and duration scores produces a CDI index value between .008 (essentially zero) and one.

These values can be measured after an event, roughly estimated by analysts with general knowledge, or predicted with more precision by somebody with detailed knowledge about how a particular organization's IT networks and procedures map onto operations. They can be deployed in debates about whether an attack against some piece of privately-owned critical infrastructure is severe enough to warrant U.S. government involvement under Presidential Policy Directive 41, an Obama-era directive outlining principles for Executive Branch response to significant cyber events in the public or private sector, and possibly even nuclear retaliation, as threatened by the Trump administration. They can be used to determine what percentage of cyber events in a data set are serious problems rather than minor annoyances. They can also help organizational leaders set priorities for prevention and risk mitigation by identifying what types of disruptive events could have the most serious consequences and should be prevented if possible; what would be the most cost-effective way to mitigate mid-level risks; and what types of cyber events can be tolerated as a cost of doing business.

Estimating the Severity of Exploitation: The Cyber Exploitation Index (CEI)

To assess the consequences of data theft, we leverage a different measure, the Cyber Exploitation Index (CEI). Total number of records compromised is not a good measure of severity, because a hacker could steal billions of records about books checked out from a library system without doing much lasting harm. The CEI score reflects both the amount and importance of three types of information: lost customer records, organizational data, and intellectual property. An event that compromised thousands of old billing records with customer addresses but no other sensitive details would receive relatively low CEI score. Events that accessed hundreds of confidential medical records or one that stole the secret recipe for a company's most popular product or a formula describing some breakthrough in weapons technology.

If multiple types of information are compromised in a single event, as occurred in the exploitative component of the 2014 Sony Pictures hack, a weighted aggregation across different types of losses can be used to assess overall severity. That event included the theft of data from the corporate e-mail server as well as documents from computers of high-ranking company officials. Organizational e-mail for the entire enterprise, including the CEO, customer details, and intellectual properties (some unreleased movies) were all stolen. Aggregating the weighted percentage of each data type lost in the compromise would reflect judgments that much of what was taken was not particularly sensitive, while the loss of a relatively small number of unreleased movies and embarrassing e-mails from senior officials was extremely costly.

Using the CDI and CEI methods to assess severity still involves some subjective judgments, but it is an improvement over the CIA grid or the PPD-41 schema because it directs users to make explicit their judgments about a standard set of considerations. Two analysts using the CDI might produce somewhat different scores, but they could quickly identify where their underlying judgments differed (for example, that impact on productivity was minimal or significant). Using a consistent method to assess severity would provide a basis for comparing how disruptive different types of cyber events at different kinds of organizations were. It could also provide a more meaningful way to assess the state of cybersecurity over time. Tracking not only how many cyber events happened in a given year, but also what their CDI or CEI scores were, would indicate whether more strategic protection decisions were mitigating risks and preventing the most severe types of cyber events, or if both the number and severity of cyber attacks were continuing to increase.

Component 4: Assessing Cyber Risks across an Organization or Sector

To make strategic decisions about risk mitigation, organizational leaders need a convenient way to compare different types of cyber risks across different parts of their organizations. Government officials with oversight responsibilities for cybersecurity across a critical infrastructure sector composed of many interconnected organizations also need some way to assess integrated risk.

To do a complete cyber risk assessment for an individual organization, the risk assessment method describe above can be repeated across all plausible scenarios and all organizational functions defined by the user (e.g. Cyber Strands). The scored scenarios associated with each organizational function can be combined in customized ways to display tables, topologic maps, or other visualizations to demonstrate the scored effect, likelihood, and accumulated risk by effect category and by organizational function.

Risk maps, called “Organizational Risk Fabrics” in CISSM framework terminology, can be generated to help technical staff communicate with organizational leadership by highlighting what type of effect poses the greatest threat to which part of their organization.

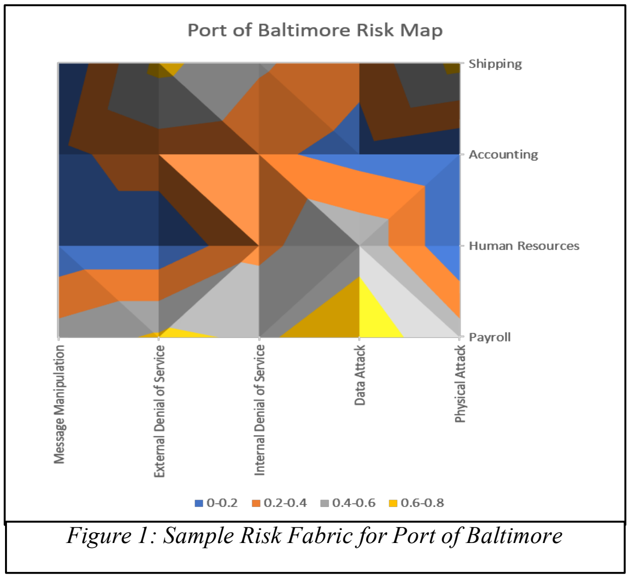

Figure 1: Sample Risk Fabric for Port of Baltimore

Figure 1 provides an example of a Risk Fabric for a set of hypothetical scenarios and scored effects against a private terminal operator at the Port of Baltimore. External denial of service attacks against the shipping and payroll groups are deemed the highest risk along with physical attacks against the shipping group and data attacks against payroll.

Organizational risks can be aggregated into integrated risk maps by combining individual cyber strands from differing organizations together into its own Risk Fabric. Breaking down specific sections of networks into functional services and then combining them together into a single view enables policy makers to maintain an integrated view of risk across multiple, independent services. These Risk Fabrics can be created for specific organizations to visualize integrated risk internally, or assembled to visualize risk for functions across organizations, but must work together to provide a public service. For example, the ability to produce and deliver electricity requires multiple organizations, and their networks, to work together in the production, transmission, and delivery of energy to residences and businesses. Only portions of the networks for each organization are required as part of this broader critical infrastructure service (e.g. payroll computers are not used to manage electrical flow). Therefore, the portions of each network that are involved can be broken down into their own cyber strands, scenarios run against each, and then combined in a single integrated Risk Fabric visualization.

This can be useful for leaders of complex organizations, like government agencies, multinational businesses, and large research universities, who might want the heads of major subordinate organizations do their own cyber risk analysis that would feed into an integrated risk analysis for the overarching organization. It can provide a way to satisfy the NIST Standard Framework principle to consider risk integrated across firms in a supply chain or some other type of interdependent relationship. And, it can provide an integrated view for policy makers who need to have a perspective across a range of independent but ultimately integrated private services.

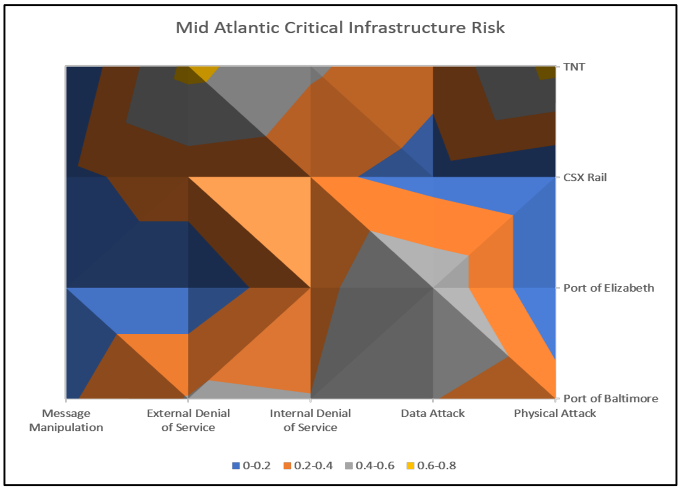

Figure 2: Sample Risk Fabric for Mid-Atlantic Transportation Risk

Figure 2 provides a sample sector-wide Risk Fabric combining scores for four organizations that are part of a regional transportation critical infrastructure sector. The greatest risk in this Risk Fabric is with freight forwarder TNT for both External Denial of Service and Physical attacks, as the underlying scenarios, likelihood estimation, and severity calculations highlight the greatest risk at an aggregated level. Message Manipulation attacks remain a relatively low risk across all services (Cyber Strands) in the fabric. This figure highlights a major strength of the CISSM framework: the ability to pull together in a modular manner the underlying impacts of cyber-attack scenarios in an integrated manner. Policy makers who are interested not only in the specific impacts to a critical organization, but also the integrated risk it poses with others as part of a broader public service they provide, gives unprecedented ability to develop customized risk topologies across a range of integrated entities.

Conclusion, Implications, and Future Research

The CISSM framework is designed to be leveraged to assess an assortment of cybersecurity risks at the device, network segment, organization, or sector level. Its principal components—a taxonomy of effect, a method for associating IT assets with organizational functions, two scoring algorithms, and a way to conceptualize risk across complex systems—can be mixed and matched to provide insight for IT experts, organizational leaders, and policy makers. Greater nuance and precision in our discussion of cyber threat and risk helps technical staff, leadership, and public officials communicate more effectively and identify the threats that move beyond a private problem and generate a more general public concern.

The components of the CISSM framework developed to date can help anybody with responsibility for cybersecurity do risk assessment for their organization or critical infrastructure sector in a way that satisfies the principles in the NIST standard framework by focusing on the primary effects of attack scenarios. CISSM is currently developing techniques for utilizing the building blocks develop in this framework to assess the secondary effects a particular attack may generate for an organization and the second order effects on an ecosystem of which it is a part (e.g. a supply chain, city, country, or transborder region). Having a more sophisticated way to define, measure, and simulate the full range of effects that a cyber event can have on different types of stakeholders will provide valuable tools to help differentiate between cyber risks that represent a manageable private problem and those that represent justifiable public concern.

About the authors

Charles Harry is a senior leader, practitioner, and researcher with over 20 years of experience in intelligence and cyber operations. Dr. Harry is the Director of Operations at the Maryland Global Initiative in Cybersecurity (MaGIC), an Associate Research Professor in the School of Public Policy, and a Senior Research Associate at CISSM. Prior to his work with the university, Dr. Harry grew and led a $35 million dollar cybersecurity consulting organization combining analysts and developers to deliver innovative solutions to the private and public sector. His public service includes a 14-year career with the National Security Agency rising to the rank of senior technical leader (DISL). Dr. Harry holds degrees in Economics and History from the University of Colorado, and was awarded a PhD in Policy Studies from the University of Maryland. He is the recipient of the Director of National Intelligence Extraordinary Achievement Medal and the Signal Intelligence Career Achievement Medal.

Nancy Gallagher is the Director of the Center for International and Security Studies at Maryland (CISSM) and a Research Professor at the University of Maryland's School of Public Policy. She leads the Advanced Methods of Cooperative Security Program, an interdisciplinary effort initiated with John Steinbruner to address the security implications of globalization by developing more refined rules and systematic transparency arrangements to minimize the potential for deliberate or inadvertent misuse of powerful, multipurpose technologies. Her current research agenda also includes space and cybersecurity policy, US and Iranian public opinion about the nuclear deal, requirements for a large-scale global expansion of nuclear energy that does not increase risks from proliferation or terrorist access, and multi-stakeholder approaches to global governance. Before coming to the University of Maryland, Gallagher worked at the State Department and the Arms Control and Disarmament Agency. She was the Executive Director of the Clinton administration's Comprehensive Test Ban Treaty Task Force. She has a Ph.D. in Political Science from the University of Illinois and a B.A. in History from Carleton College.