Cyber Risk Scoring for Executives: An Assessment Framework

Abstract

Cybersecurity data is robust and prevalent. However, that data is rarely correlates vulnerabilities, security controls, and threats for a target environment. As cyber related threats continue to increase in prevalence and potential impact, organizations must become more agile and strategic in understanding and responding to threats. The following paper provides a framework inputs to calculate a numerical risk score and present the information in a dashboard for use by executives and security teams. The framework includes strategic business and technical considerations, along with a rudimentary equation to aggregate cybersecurity data in near-real time.

Introduction

Cybersecurity threats are prevalent and expanding in both the realization of adversarial action and impact on individuals and organizations. As stated in Verizon's 2019 Data Breach Investigations Report, "Seemingly, no matter what defensive measures security professionals put in place, attackers are able to circumvent them. No organization is too large or too small to fall victim to a data breach. No industry vertical is immune to attack. Regardless of the type or amount of your organization's data, there is someone out there who is trying to steal it."

The impact that bad actors can have on an individual or a business (large or small) could be catastrophic. According to the Ponemon Institute, the total average cost to an organization for a data breach in 2017 was $3.62 million and the average cost per individual stolen record containing sensitive and/or confidential information was $141. Combined with the increasing

prevalence of attacks, these monetary values illustrate the serious nature and widespread effect of cybersecurity incidents, targeting not only organizations and businesses, but also individuals. While organizations that retain sensitive or personal information are ultimately responsible for its safe handling, a breach of confidentiality also results in impact to the individual to which the information pertains. This has led to numerous instances of businesses paying for identity protection services for individuals who may have been impacted by breaches targeting data for which the organization was the responsible custodian.

I propose that these facts raise a key question: how do we better protect ourselves and our businesses? Protections that need to be aimed at limiting exposure to loss of financial well-being, public perception, time spent restoring personal accounts, and reducing the need for additional insurance vehicles (like cyber-insurance or identity theft protection). Aside from investing more into ongoing cybersecurity programs that attempt to apply a myriad of holistic protections in the rat race with adversaries, what innovative approach can change the way that we view the problem? How do we effectively analyze the droves of data that we have about our IT environments and correlate it with available breach reports, threat intelligence and studies that analyze adversarial actions taken against an identified target?

The improvements and investments in cybersecurity analytics and enabling platforms like security information and event management (SIEM) tools and threat intelligence sharing platforms have empowered organizations to detect and respond to breaches (Ponemon). The number of days it took organizations to identify a data breach decreased from 201 days in 2016 to 191 days in 2017 and the average number of days to contain the breach decreased from 70 days to 66 days (Ponemon). While these numbers are trending in the right direction, it takes

257 days (over 8 months) on average to detect and contain a data breach (Ponemon). That leaves an attacker with 8 months to understand an organization's IT infrastructure, the cybersecurity teams, security gaps that could be exploited, and activate processes that silently exfiltrate identified sensitive and confidential data.

However, I pose that significant opportunities exist to aggregate information in a targeted, actionable manner that empowers organizations to manage security protections and investments for their most critical assets or attack vectors. I do not believe it is plausible for organizations to solve cybersecurity problems by making uninformed, blanket investments into cybersecurity programs that place the same value on all threats and assets. The IT and cyber related data that many organizations have available for consumption is continually being enriched and can be better leveraged. This will assist boards, C-suite executives, and senior level management in understanding where the greatest cybersecurity risk and potential impact resides for their organization and will inform sound investment strategies and goal setting. In this paper I will present the factors and key considerations to be made in developing an algorithm to quantify risk into a numerical value for IT systems, business units, and organizations.

Approach

To effectively analyze this issue, I leverage existing industry recognized frameworks to articulate and approach the facets of cybersecurity. This organizes the analysis into a pre-defined and well-known structure and lexicon, limiting the need to expound upon processes and terminology. Examples of these frameworks include the National Institute of Standards and

Technology (NIST) Cybersecurity Framework (CSF), the International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC) 27000 series, the Center for Internet Security (CIS) Critical Security Controls (CSC), guidance from the Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency (CISA), and the Open Web Application Security Project (OWASP), among others. The approach utilizes these frameworks as a guideline to integrate cybersecurity best practices, attack vectors, and available organizational cyber data in an actionable way that may be tailored to an individual business or industry. Evaluating these data points in the context relevant threat intelligence information will enable the creation of a numerical risk score. This score, along with supporting evidence, should then be presented in a dashboard that illustrates the posture of the organization, individual systems and associated assets, and quantifies the likelihood and potential impact of a cyber incident.

This framework expects that the calculated score is available in near-real time. This is a significant transformation from the traditional approach to organizational risk-assessments, which provide a point in time perspective at infrequent intervals. This limited assessment may occur at a high-risk or low-risk point in time for a system or for the organization (sample below), skewing the analysis and resulting impact of investments.

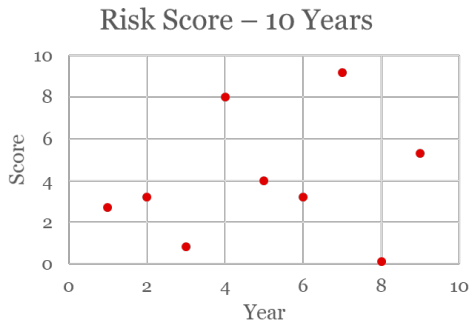

In contrast, recurring scores available in near-real time (below) empower organizations to view their risk posture over time with less overhead, providing a more accurate depiction of the environment and driving more impactful decision-making. These additional data points can be used to calculate an average score for a target asset, system, or organizational unity, and can mitigate against outlier fluctuations in risk.

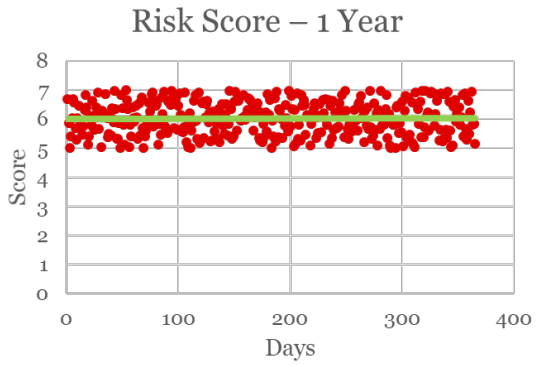

Scoring Calculation Assumptions

This framework does not provide a specific means for calculating risk scores but requires that some basic assumptions be made about how those calculations take place. First, the score at the organizational level is an aggregation of scores assessed at the system and asset levels. This creates a tiered approach, where the summation of identified cybersecurity metrics (vulnerabilities, configuration deficiencies, etc.) and impact define the risk scores for individual assets, the aggregation of asset scores define the risk of systems, and the aggregation of system scores result in the organizational risk score. In addition, the framework assumes that each individual metric, data point, or score can be numerically driven to a higher or lower risk profile using a coefficient. The coefficient should be designed to fluctuate based on the incorporation of guidance described in the framework and reflect business specific and real-world threats.

Leveraging Frameworks

In conducting research and considering past experiences, it is important to both leverage a standardized framework to structure the analysis as well as to understand the nature, context, and maturity level of the organization being evaluated. In the analysis presented I primarily use the NIST CSF to structure and analyze cybersecurity components. This analysis is broken down into six primary categories: understanding the framework, calculation assumptions, setting the stage, data aggregation, evaluation and scoring, and developing and using the dashboard.

The NIST CSF provides a holistic and technology agnostic approach to establishing a cybersecurity program, broken down into five functions: Identify, Protect, Detect, Respond, and Recover (CSF). These functions are further broken down into categories and subcategories and mapped to other cybersecurity frameworks (ISO 27000, CIS CSC, COBIT 5, NIST 800-53, and ISA 62443). For the purposes of this paper, focus will be placed on categories outlined in the Identify, Protect and Detect phases of the NIST CSF. However, it is important to note that information gained from the Respond and Recover phases would be beneficial to inform risk scores and investment activities. When incorporated, the resulting fluctuations in scores will be specific to events realized and will be left for future research and development of this framework.

Setting the Stage

Business Drivers and Alignment

This is perhaps the most important aspect in enabling actionable and effective outcomes and can often be overlooked or undervalued. It is imperative that organizations understand, dedicate effort for, and clarify what is critically important for their operation and business model. Beyond an internally focused self-assessment, the organization must consider perceptions of the media and general public, as these perceptions may influence adversarial targets and tactics, techniques, and protocols used. In turn, the risk calculation can leverage assessment information to make assumptions about the attack surface and generate baseline impact coefficients. This assessment process should be iterative and recurring to maintain a current view of the organization and its threats.

Assessing the organization begins with understanding the business environment. In reference to the NIST CSF, this is part of the Identify function in the Business Environment category. Key differentiators exist between industry verticals that require special consideration in assessing risk. These focus areas can be easily incorporated into the scoring calculation if effectively pre-defined by the business. A non-exhaustive list of questions to be answered and the potential impact on assessing risk is included below.

Question	Impact
What industry vertical is the organization most closely aligned to?	Industries have different targets for attackers that are working to maximize impact while minimizing work. Some examples include (Koegler): • **Healthcare*** – Personally identifiable information (PII), personal health information (PHI), electronic medical records (EMRs), etc. • **Finance** – Financial transaction availability, third and fourth-party providers with access to financial data connections. • **Energy and Utilities** – Grid management, Internet of Things (IoT), system availability, etc.
Is the organization aligned to DHS's critical infrastructure sectors? If so, which sector?	Alignment to critical infrastructure may indicate an increased risk of attack for certain organizations, especially in times of political unrest or known attacks from advanced persistent threat (APT) groups.
What are the business drivers for the organization?	These questions are all designed to assist in the identification of IT systems containing or associated with impactful business drivers or sensitive information. These questions may be answered as part of operational planning activities, but they play a vital role in establishing a baseline understanding of which systems bad actors will seek to access or disrupt. This highlights the importance of coordination between IT operations, IT security, and functional business units.
What are the organizations competitive advantages?	
What intellectual property (IP) does the organization have? How is IP defined and identified?	
What types of sensitive data are collected and retained (PII, PHI, financial, IP, national security, etc.)?	
What is the level of dependence on the availability of IT systems?	Seemingly generic, this question is intended to determine the importance of integrity and availability of systems. While confidentiality of data is often a point of focus for cybersecurity teams, the integrity of information and availability of key systems may wreak greater havoc if impacted.
What applicable legal requirements exist for the protection of data, availability of systems, or reporting of loss or theft through cyber-crimes?	Used to identify and assess additional layers of financial or reputational risk to systems and/or business units.

*"Your medical information is worth 10 times more than your credit card number on the black market." (Reuters)

These questions establish a baseline to cross-reference business impact against a list of IT systems, including the function of the system and data it collects and/or stores. Alignment of IT and business drivers enable the prioritization of systems to generate a targeted view of risk, based on potential impact to the organization. Impact should be considered (where applicable) in terms of productivity, financial (including real cost of loss, estimated cost of future loss, response efforts, re-building, and customer protection), national security, public health, loss of competitive advantage, and reputation. Each of the above considerations, the potential impact, risk tolerance of the organization, as well as a standardized set of data types and potential risks for the industry should be applied to each system.

Cyber Program Maturity

In addition to a general assessment of the business and IT environment, the organization must have a mature cybersecurity program for the risk scoring to be actionable. The maturity level of the cybersecurity program can be adequately assessed utilizing the NIST CSF. Using this framework, organizations will have the greatest benefit if at least a "Tier 3: Repeatable" (CSF, p. 10) maturity level has been achieved within the business unit or system(s) being evaluated. As described in the framework, this tier includes: 1) risk management policies that are formally approved and practices regularly updated, 2) an organization-wide approach to manage cybersecurity risk, and 3) organizational understanding of its role in the larger ecosystem and the broader context of risk. While it is ideal to have these concepts applied across the organization, it is minimally required for the target of the assessment, if limited to a singular system or unit of the organization. For example, to effectively conduct this assessment, it is imperative for organizations to have a robust and mature system inventory

and asset management process, else it would not be possible to effectively determine what to evaluate and protect.

The assessment and maturity evaluation should include a detailed technical understanding of the IT environment landscape, including an alignment between IT systems and infrastructure and key systems identified for storage of sensitive data, intellectual property, or core business operations. This mapping must align with the drivers defined in the evaluation of the business. Data classification should be completed and documented to define sensitive data types for the organization (inclusive of common data types identified for an industry) and appropriately align data for each system to be scored. This exercise is intended to ensure that the information pertinent to the business, along with the level of impact, has been applied to the inventory of systems.

While the technical environment is the central focus of the framework, people in the organization being assessed are crucial in maintaining security (Ramachandran). A mature training and awareness program is required to influence culture, protect against threats early in the system development lifecycle, and engage the organization in identifying and reporting potential cybersecurity events. In assessing the maturity of this capability, organizations should consider the ability to quickly and effectively communicate to users, the veracity of trainings provided, and performance evaluations or practical assessments of cybersecurity knowledge. This will impact the risk scoring by either increasing or decreasing the overall baseline score for the organization.

Combined, each of the factors in 'Setting the Stage' will adjust the levers used in determining the potential impact coefficient applied to the risk score. This is a manual process that will establish a baseline of the weighting coefficients which will be further tailored by the compensating controls and exploitable risk in the technical environment.

Data Aggregation

Perhaps the most important factor in the cybersecurity maturity level of the organization is the ability to effectively collect, retain, and make available for analysis relevant cybersecurity data. This data must not only be comprehensive in its coverage of assets in the IT environment, but the organization must have a high-degree of confidence in the validity of the data gathered. To enable near real-time calculation of risk scores, organizations should be able to obtain robust data about each of the following: system and asset inventory, audit and assessment results, configuration and compliance scans, anti-malware logs, vulnerability scans (servers, workstations, databases, web applications, mobile devices, etc.), software use and code execution, data movement and encryption, network traffic inspection, network baseline use, threat intelligence, user behavior analytics and access control information, event log aggregation, change logs, incident data, code scanning, and penetration test findings. This list is non-exhaustive but is intended to capture a cross section of cybersecurity related data to make an effective determination of risk. It is also important that the list adapt to new technologies and applicable data sources.

In addition to information collected about the potential vulnerabilities or anomalous behavior that exists within an environment, organizations should also capture compensating

controls. The representation of defensive measures applied should provide sufficient information for a data aggregation system or the scoring framework to determine the scope of risk to be reduced and the types of vulnerabilities that may be mitigated. The application of broad security measures such as Internet Protocol Security (IPSec), Domain Name System Security (DNSSec), or zero-trust models must be accounted for. Organizations should not only understand what the potential attack surface looks like, but also take credit for and incorporate the actions taken to reduce the likelihood of an attack. The incorporation of compensating controls into the framework will also enable the data-driven realization of benefits of cybersecurity investments.

Adversaries are work-averse and will employ attack mechanisms that have the potential to impact the greatest number of vulnerable users (Allodi et. al.). The caveat to this statement is represented by advanced persistent threat (APT) actors that have the time and resources at their disposal to utilize non-traditional means to attack a subject. In either case, actions that can be taken to increase the difficulty of an attacker will act as a deterrent and decrease the likelihood of a cybersecurity event.

The vulnerability and security control data that is collected must be aggregated in a way that each data point can be associated with an asset or system, whichever is more appropriate. This is imperative because it not only enables the base calculations that lead to an aggregate risk score, but it allows for the information to be represented at multiple levels of technical complexity in the dashboard. While the aggregate score is designed for the executive level, security teams will be concerned with the specific vulnerabilities or configurations that can be mitigated or altered to lower the score. As will be outlined below, the resulting scores and data

points would be best depicted in a singular source (dashboard) that allows executives to view aggregate information while enabling security teams to obtain granular details.

In an ideal state, organizations would have a high degree of confidence that the data that is being collected is accurate and representative of the entire environment, however this is generally not the case. For organizations at the earlier stages of maturity, the following critical elements should be gathered for scoring purposes: system and asset inventory, audit and assessment results (where applicable), configuration and compliance scans, anti-malware logs, vulnerability scans (servers, workstations, databases, web applications, mobile devices, etc.), network traffic alerting, and security log aggregation.

Evaluation and Scoring

Following the establishment of organizational and impact coefficient baselines, along with the model of for collecting and assigning pertinent data points to assets, the information must be analyzed to determine a risk score. This step will take the organizational analysis information along with the real-time cyber data and compensating controls and evaluate that information against known threats to generate a risk score for the target object (asset, system, business unit, etc.). As previously outlined in the calculation assumptions, the risk scores can then be aggregated to generate a score for the enterprise.

There are a vast number of nuances that could be considered when converting cyber data into risk scores. Vulnerability scan results are the most likely data point to be impacted by factors other than what is reported by detection tools. While this is by design, as the Common Vulnerability Scoring System (CVSS) is formulated to be applicable across all environments,

much of the context of threat vectors is lost in failing to adapt the impact to the target environment. For example, compensating controls may be implemented that protect a vector of an attack related to a vulnerability without impacting the vulnerability signature used in detection. The detection tool will still indicate to the organization that a vulnerability exists and action must be taken.

Keeping this in mind, CVSS scoring should be used as a baseline to deduce risk to the organization. CVSS scoring information should then be evaluated alongside compensating controls that reduce the possibility of exploit as well as threat intelligence that may indicate a higher level of risk based upon active adversarial campaigns. In addition, greater risk should be accounted for vulnerabilities that have an exploit contained in the most commonly available and used exploit kits. In 2011, Guido illustrated that the top four exploit kits were responsible for 85% of the volume of malicious URLs, while the top 10 kits were responsible for 99%. As this study was conducted in 2011, the specific kits may not represent the same share of the market today. However, the representation that relatively few exploit kits control the market is indicative of the availability of certain kits to adversaries. Some additional factors for consideration in scoring are represented in the table below:

Subject	Consideration
Asset Criticality	Potential risks are not equal for all systems or assets. Information from the system evaluation should be used to determine criticality to business operations, data protected, whether it is part of a development, staging/testing, or production environment or is externally (public/internet) facing.
Web Application Vulnerability	Similar to the consideration for vulnerabilities that exist as part of common exploit kits, web application vulnerabilities that are associated with the Open Web Application Security Project (OWASP) Top Ten list should be

Subject	Consideration
	granted a higher risk score. This is due to the prevalence of information regarding these attack types, the ease of development of exploits, and the potential impact that could be inflicted on an organization.
Threat Intelligence	Threat intelligence feeds can be used to inform attack vectors, adversaries, and industries that are at a higher risk of impact. Less mature implementations of this factor will affect the risk scoring for the enterprise, while more detailed information could be used to influence scoring for specific vulnerable systems or those that contain at-risk data types.
Network Use	The baseline of network use should be captured as a data point in order to identify potentially anomalous behavior that is indicative of adversarial action, resulting in an increase to the risk score. However, organizations must keep in mind that anomalous behavior may have already been occurring at the time of baselining and place additional scrutiny on what is deemed as authorized activity.
System Interconnections	Special consideration should be given to assets or systems that have a direct connection to those that are indicated as higher impact for the organization. This will increase risk scores for these systems to account for the relative ease of lateral movement achieved by adversaries.
Least Privilege	While least privilege should be included as a compensating control, it is important to note its impact to overall scoring techniques due to its ability to greatly increase the work on an attacker and hinder lateral movement as well as any access to critical systems. When applied holistically, least privilege can be the difference between attacker access to an individual computer and access to all critical systems.
Audits, Assessments and Penetration Tests	Assessment findings that are related to controls listed in the Center for Internet Security (CIS) Critical Security Controls (CSC) should be considered with higher risk to the organization. This evaluation should be in context with or in addition to the risk and impact level identified by the audit methodology being used.

A rudimentary application of this framework into a calculation could be demonstrated as the summation of all data inputs, with impact coefficients utilized to increase or decrease the potential risk factor of a point. For example, a specific asset may be evaluated as follows, where the impact coefficient is represented as (IC):

(IC^1 Asset Criticality) * [(IC^2) * Total High Vulnerabilities + (IC^3) * Total Moderate Vulnerabilities + (IC^4) * Total Low Vulnerabilities + (IC^5) * Web App Vulnerabilities + (IC^6) * Configuration Deviations …]

To be actionable, each calculated data point must be evaluated and scored in a consistent manner and aggregated for each system, followed by the aggregation of all systems to result in a score for the organization. This will lead to a high-level depiction of overall cybersecurity risk that is rooted in verifiable data and real-world threats.

Representation of Scores and Actionable Intelligence

The scoring framework is the focus of this paper, however the importance of the representation of results to the intended audience cannot be overlooked. The impact and importance of pertinent and valuable information can easily be misplaced if it is not presented at the appropriate level of detail and using common language. To that end, once scores have been calculated an aggregate view should be provided through a web-based dashboard. This can either utilize a commercial off the shelf (COTS) product for data visualization or be incorporated into the data aggregation tool itself, such as the SIEM or governance risk and compliance (GRC) application. The default view should be focused on the executive audience and include the following (sample mock-up below):

- The overall enterprise risk score at the top, color-coded for the current state;
- A list of systems with brief description of function and the risk score for each system;

- The user should have the ability to customize the view presented, but default should show critical systems and/or those with high risk listed first;
- System names and assets listed should allow users to navigate to details about what is contributing to the risk score;
- Information on each page should be well organized and structured to avoid over-burdening the user; and
- All information must be presented using easily understandable (non-technical) language.

Organization	Score	Trend	Average
Rand Corp	796		847

System Name	Description	Impact	Score	Primary Risk Factors	Trend
Critical System 1	Runs the business	High	674	VULN, ENC	
Critical System 2	Operates IT	High	563	NET, TH.INT.	
Intermediate 1	Connects important functions	Moderate	987	LOGS, ENC	
Critical System 3	Stores all the data	High	32	DB VULN, COMPLIANCE	
Intermediate 2	Admin functions	Moderate	234	ACCESS, VULN	
Low 1	Plays games	Low	463	ACCESS	

Conclusion & Future Work

The framework outlined here is intended to enable the effective assessment of cybersecurity risk at an organization, which will in turn enable strong executive decision making and informed investments. As is suggested by Allodi et. al., increasing the work on an attacker can be an effective means in protecting against adversarial activity. In elevating the exposure of

cybersecurity risks to the executive level in an effective way, organizations will be empowered to define strategic cybersecurity objectives, reduce the time spent analyzing and defining budget requests, and fund initiatives that effectively reduce risk.

The natural follow-on to this work is to establish a concrete mechanism for the calculation of the risk score, as well as to define baseline impact coefficients for individual industries. Following its creation, a pilot and evaluation of the scoring mechanism in real-world scenarios will provide evidence of its efficacy along with identified adjustments to baselines for specific organizations and industries. It is also important to note that in the ever-evolving world of cybersecurity, the scoring mechanism must be flexible enough to incorporate new data points, risk factors, device types, or other factors that would be relevant to an organization's understanding of cybersecurity risk.

Works Cited

1. "2019 Data Breach Investigations Report: Executive Summary." *Verizon*, 2019, enterprise.verizon.com/resources/executivebriefs/2019-dbir-executive-brief.pdf. Accessed 23 February 2020.
2. "2017 Cost of Data Breach Study: Global Overview." *Ponemon Institute*, June 2017, ibm.com/downloads/cas/ZYKLN2E3. Accessed 23 February 2020.
3. "Cybersecurity Framework." *National Institute of Standards and Technology (NIST),* nist.gov/cyberframework. Accessed 23 February 2020.
4. "ISO/IEC 27001 Information Security Management." *International Organization for Standardization (ISO)*, iso.org/isoiec-27001-information-security.html. Accessed 23 February 2020.
5. "Center for Internet Security Critical Security Controls." *Center for Internet Security (CIS)*, cisecurity.org/wp-content/uploads/2017/03/Poster_Winter2016_CSCs.pdf. Accessed 23 February 2020.
6. "Cybersecurity & Infrastructure Security Agency." *Department of Homeland Security (DHS) Cybersecurity & Infrastructure Security Agency (CISA)*, cisa.gov/cybersecurity-division. Accessed 23 February 2020.
7. "Open Web Application Security Project." *Open Web Application Security Project*, owasp.org/. Accessed 23 February 2020.
8. "New ISA/IEC 62443 standard specifies security capabilities for control system components." *The International Society of Automation*, isa.org/intech/201810standards/. Accessed 23 February 2020.

9. Koegler, Scott. “Vulnerabilities Across Verticals: Overcoming Industry-Specific Security Challenges.” *SecurityIntelligence*, 26 September 2017, securityintelligence.com/vulnerabilities-across-verticals-overcoming-industry-specific-security-challenges/. Accessed 23 February 2020.
10. “Critical Infrastructure Sectors.” *Department of Homeland Security (DHS) Cybersecurity & Infrastructure Security Agency (CISA)*, cisa.gov/critical-infrastructure-sectors. Accessed 23 February 2020.
11. Humer, Caroline and Finkle, Jim. “Your medical record is worth more to hackers than your credit card.” *Reuters*, 24 September 2014, reuters.com/article/us-cybersecurity-hospitals/your-medical-record-is-worth-more-to-hackers-than-your-credit-card-idUSKCN0HJ21I20140924. Accessed 23 February 2020.
12. Ramachandran, Ramesh. “The Importance of Training: Cybersecurity Awareness like a Human Firewall.” *Entrepreneur*, 15 October 2019, entrepreneur.com/article/340838. Accessed 23 February 2020.
13. Allodi, Luca, Massacci, Fabio, and Williams, Julian. “The Work-Averse Cyber Attacker Model: Theory and Evidence From Two Million Attack Signatures.” *The Work-Averse Cyber Attacker Model: Theory and Evidence From Two Million Attack Signatures*, 29 May 2017, cs.brown.edu/people/jsavage/EMCS2600Readings/Module2/2017_WEIS_WorkAverseCyberAttackModel.pdf. Accessed 23 February 2020.

14. "Common Vulnerability Scoring System v3.1: Specification Document." *Forum of Incident Response and Security Teams*, first.org/cvss/v3.1/specification-document. Accessed 23 February 2020.

15. Guido, Dan. "A Case Study of Intelligence-Driven Defense." *IEEE Security & Privacy*, Dai Zovi, Dino, and Sotirov, Alexander, IEEE Computer And Reliability Societies, November/December 2011, cs.brown.edu/people/jsavage/EMCS2600Readings/Module1/2011_IEEE_SecurityPrivacy_IntelligentDefense_Guido.pdf. Accessed 23 February 2020.