



Prepared for:
Department of Homeland Security

Cyber Risk Metrics Survey, Assessment, and Implementation Plan

May 11, 2018

Authors:

Nathan Jones
Brian Tivnan

The Homeland Security Systems Engineering and Development Institute (HSSEDI)[™]
Operated by The MITRE Corporation

Approved for Public Release; Distribution Unlimited.
Case Number 18-1246 / DHS reference number 16-J-00184-05

This document is a product of the Homeland Security Systems Engineering and Development Institute (HSSEDI[™]).

Homeland Security Systems Engineering & Development Institute

The Homeland Security Act of 2002 (Section 305 of PL 107-296, as codified in 6 U.S.C. 185), herein referred to as the “Act,” authorizes the Secretary of the Department of Homeland Security (DHS), acting through the Under Secretary for Science and Technology, to establish one or more federally funded research and development centers (FFRDCs) to provide independent analysis of homeland security issues. MITRE Corp. operates the Homeland Security Systems Engineering and Development Institute (HSSEDI) as an FFRDC for DHS under contract HSHQDC-14-D-00006.

The HSSEDI FFRDC provides the government with the necessary systems engineering and development expertise to conduct complex acquisition planning and development; concept exploration, experimentation and evaluation; information technology, communications and cyber security processes, standards, methodologies and protocols; systems architecture and integration; quality and performance review, best practices and performance measures and metrics; and, independent test and evaluation activities. The HSSEDI FFRDC also works with and supports other federal, state, local, tribal, public and private sector organizations that make up the homeland security enterprise. The HSSEDI FFRDC’s research is undertaken by mutual consent with DHS and is organized as a set of discrete tasks. This report presents the results of research and analysis conducted under:

HSHQDC-16-J-00184

Next Generation Cyber Infrastructure (NGCI) Apex Cyber Risk Metrics and Threat Model Assessment

This HSSEDI task order is to enable DHS Science and Technology Directorate (S&T) to facilitate improvement of cybersecurity within the Financial Services Sector (FSS). To support NGCI Apex use cases and provide a common frame of reference for community interaction to supplement institution-specific threat models, HSSEDI developed an integrated suite of threat models identifying attacker methods from the level of a single FSS institution up to FSS systems-of-systems, and a corresponding cyber wargaming framework linking technical and business views. HSSEDI assessed risk metrics and risk assessment frameworks, provided recommendations toward development of scalable cybersecurity risk metrics to meet the needs of the NGCI Apex program, and developed representations depicting the interdependencies and data flows within the FSS.

The results presented in this report do not necessarily reflect official DHS opinion or policy.

For more information about this publication contact:

Homeland Security Systems Engineering & Development Institute

The MITRE Corporation

7515 Colshire Drive

McLean, VA 22102

Email: HSSEDI_info@mitre.org

<http://www.mitre.org/HSSEDI>

Abstract

The Homeland Security Systems Engineering and Development Institute (HSSEDI) assists the Department of Homeland Security (DHS) Science and Technology Directorate (S&T) in the execution of the Next Generation Cyber Infrastructure (NGCI) Apex program. This report identifies existing metrics and surveys relating to cybersecurity, and provides an Implementation Plan for an NGCI Confidence Survey. It culminates with three main recommendations: 1) the NGCI program should develop a scalable framework for cybersecurity risk metrics; 2) the NGCI program should administer a Confidence Survey to the Cyber Apex Review Team (CART); and 3) the Confidence Survey should serve as an initial step on the trajectory to a scalable cybersecurity risk metric.

Key Words

1. Next Generation Cyber Infrastructure (NGCI) Apex program
2. Critical Infrastructures
3. Cybersecurity Risk Metrics
4. Cyber Threat Models
5. Confidence Surveys

This page intentionally left blank

Table of Contents

1	Project Overview.....	1
1.1	Cyber Risk Metrics Task	1
1.2	Overview of Recommendations	2
2	Background.....	3
2.1	Definition of Risk	3
2.2	Overview of Risk Metrics	4
2.3	Risk Metrics in the Financial Services Sector (FSS)	6
2.3.1	Government.....	6
2.3.2	Industry.....	7
2.3.3	Academia.....	8
3	Cybersecurity Risk Metrics	10
3.1	Risk Metric Development Frameworks	10
3.1.1	Risk Models	10
3.1.2	Cybersecurity Investment Framework	11
3.1.3	Framework-Based Tool	14
3.2	Quantitative Approaches	14
3.2.1	CVaR	14
3.2.2	FAIR.....	15
3.2.3	Industry Developed Metrics	16
3.3	Red Teaming.....	17
4	Cybersecurity Risk Metrics Survey of FSS	18
4.1	Process.....	18
4.2	Data.....	19
4.3	Reporting.....	19
4.4	Testing & Validation	20
5	Synthesis of Cybersecurity Risk Metrics Findings	21
6	Analysis and Assessment.....	22
6.1	Assessment Methodology	22
6.2	Families of Cybersecurity Metrics.....	22
6.3	Cybersecurity Metric Assessment Criteria.....	23
6.4	Cybersecurity Metric Assessment Findings	24
6.4.1	Quantitative Risk Metrics	26

6.4.2	Quantitative Non-Risk Metrics	27
6.4.3	Qualitative Risk Metrics.....	28
6.4.4	Qualitative Non-Risk Metrics.....	28
6.5	Synthesis for the NGCI Apex Program	28
7	Quantitative Approach to Cybersecurity Risk Metrics.....	30
7.1	Argument in Favor of a Quantitative Approach to a Cybersecurity Risk Metrics Program ...	30
7.1.1	Transparency	30
7.1.2	Repeatability.....	30
7.1.3	Reproducibility.....	30
7.1.4	Scalability	30
7.2	Argument Against a Quantitative Approach to a Cybersecurity Risk Metrics Program.....	31
7.2.1	Reliance on Input Data	31
7.2.2	Confusion with Precision	31
7.2.3	Mitigations	31
7.3	Candidate Quantitative Frameworks for Implementation	32
7.3.1	FAIR.....	32
7.3.2	Hubbard and Seiersen Approach	33
7.3.3	Comparison Between These Two Approaches	34
7.4	Cybersecurity Situational Awareness Through a Confidence Survey	34
8	Introduction to Confidence Survey Exemplars	36
8.1	Index of Cyber Security.....	36
8.2	Cost of Data Breach Study: United States	37
8.3	Global Cybersecurity Status Report	38
8.4	U.S. State of Cybercrime Survey	38
9	Assessment of Confidence Survey Exemplars.....	39
9.1	Confidence Survey Exemplar Assessment Methodology	39
9.2	Confidence Survey Exemplar Assessment Criteria.....	40
9.2.1	Criterion One	41
9.2.2	Criterion Two.....	41
9.2.3	Criterion Three	41
9.2.4	Criterion Four	42
9.2.5	Criterion Five.....	42
9.2.6	Criterion Six.....	42

9.3	Confidence Survey Exemplar Assessment Results	43
9.3.1	Assessment for Index of Cyber Security	44
9.3.2	Assessment for Global Cybersecurity Status Report.....	44
9.3.3	Assessment for Cost of Data Breach Study: United States	44
9.3.4	Assessment for US State of Cybercrime Survey.....	45
10	Implementation Plan for an NGCI Confidence Survey	46
10.1	Benefits of Implementing ICS.....	46
10.1.1	Administering ICS to CART Members	47
10.1.2	Quantifying the Perception of Cybersecurity Risk Among CART Members	47
10.2	Expanding ICS to Quantify Impact of Past Investments	48
10.3	Expanding ICS to Quantify Demand for Future Investments.....	49
10.4	Expanding ICS to Inform Cybersecurity Risk Modeling	49
10.4.1	Need for an Improved Cybersecurity Risk Metric	50
10.4.2	Ability to Evaluate Different Investment Alternatives	50
10.4.3	Furthering Development of Cybersecurity Risk Modeling Framework	51
11	Conclusions and Recommendations.....	52
Appendix A	Methodology for Survey of Metrics	53
List of Acronyms	55
List of References	57

List of Figures

Figure 1. Risk Is Experienced at Different Scopes [Adapted from Bodeau 2014]	4
Figure 2. Spectrum of Risk Metric Approaches	5
Figure 3. Cybersecurity Economic Framework [Adapted from AFCEA 2013]	12
Figure 4. Extended Cybersecurity Framework [Adapted from AFCEA 2014]	13
Figure 5. Metrics for Enhanced Cybersecurity Framework [Adapted from AFCEA 2014]	14
Figure 6. Cybersecurity Metric Assessment Methodology	23
Figure 7. Standard Risk Matrix	32

List of Tables

Table 1. Cybersecurity Metric Assessment Color Key	25
Table 2. Quantitative Risk Metrics Assessment.....	26
Table 3. Quantitative Non-Risk Metrics Assessment	27
Table 4. Qualitative Risk Metrics Assessment	28
Table 5. Qualitative Non-Risk Metrics Assessment.....	28
Table 6. Cybersecurity Surveys Included in Review	36
Table 7. Evaluated Result Publications	40
Table 8. Cybersecurity Survey Assessment Matrix	43

1 Project Overview

The Next Generation Cyber Infrastructure (NGCI) Apex Program seeks to accelerate the adoption of effective information technology (IT) security risk mitigating cyber technologies by the Financial Services Sector (FSS), one of the most technologically advanced critical infrastructures. By doing so, its goals are to 1) increase financial sector-wide situational understanding of evolving IT security risk and the technology associated with that risk; 2) improve the ability to understand and link compromises in the underlying cyber infrastructure to sub-sector operations; 3) enable greater information flows across sub-sectors; and 4) enable financial sector institutions to detect and neutralize adversaries more quickly and effectively than is possible today.

Understanding cybersecurity risk requires the adoption of some form of cybersecurity risk metrics. Metrics are driven by various types of risk assessments, which in turn require a credible model of threats as an essential input. The specific objective of the Cyber Risk Metrics Task is to identify cybersecurity risk metrics that could be candidates for use in the NGCI Apex program. The specific objective of the companion Threat Models Survey and Assessment is to identify threat models and frameworks that could be candidates for use in the NGCI Apex program.

1.1 Cyber Risk Metrics Task

The goal of this task is to develop cyber risk metrics that could be used to assess the impact of the NGCI program. The Department of Homeland Security (DHS) Science and Technology Directorate (S&T) expects that this work will be performed over the life of the NGCI program, and HSSEDI shall assess new technologies, tools, or techniques as determined by S&T. In the base period of this task, HSSEDI:

- Conducted a survey of metrics relevant to the NGCI program. The survey consisted of a literature review and interviews with FSS executives focusing, in part, on cybersecurity metrics used within the sector. See sections 2 through 5 and Appendix A.
- Conducted an assessment to determine the extent to which select metrics identified during the FSS executive interviews benefit their originating organization and have the potential to benefit the NGCI program. See section 6.
- Formulated a recommendation that the NGCI program should pursue the development of a quantitative cybersecurity risk metrics framework for the sector. See section 7.
- Reviewed and assessed existing surveys related to cybersecurity. See sections 8 through 9.
- Developed an Implementation Plan for the NGCI program to apply a Confidence Survey to gather feedback from the Cyber Apex Review Team (CART). See section 10.

This paper covers the activities conducted in the base period of the Cyber Risk Metrics Task. The analyses presented in this paper were conducted between May 2016 and May 2017. Generally, these analyses have not been updated to reflect new information made available after they were conducted.

1.2 Overview of Recommendations

This report identifies existing metrics and surveys relating to cybersecurity, and provides an Implementation Plan for an NGCI Confidence Survey. It culminates with three main recommendations: 1) the NGCI program should develop a scalable framework for cybersecurity risk metrics, drawing on concepts from two prominent modeling approaches; 2) the NGCI program should implement a Confidence Survey to continuously gather feedback regarding cybersecurity issues that impact the CART; and 3) the Confidence Survey should serve as an initial step on the trajectory to a scalable cybersecurity risk metric which meets the needs of the NGCI program.

2 Background

The mitigation of risk is the foundation of any institution in the FSS; indeed, some institutions will go so far as to say that risk mitigation is the single most important, competitive advantage of any institution in the FSS. It then stands to reason that risk measurement serves as the cornerstone of any effective risk mitigation plan. In this section, HSSEDI defines risk and risk metrics from the perspectives of three distinct classes of stakeholders: government, industry and academia.

2.1 Definition of Risk

A prevailing treatment of risk in the financial literature comes from Holton [Holton 2004]. Risk entails two essential components: uncertainty and exposure. Below, we will elaborate on uncertainty and exposure, but we first require an intermediate step of introducing another component – proposition.

A proposition is a statement or assertion that expresses an observation, judgment or opinion. For example, the observation that it is raining serves as a proposition which can be either true or false, but not both. Propositions are tools commonly found in varied domains from philosophical debates to scientific analyses to our observations in everyday life.

Returning now to our definition of risk, we stated above that risk entails two essential components: uncertainty and exposure. Uncertainty reflects the state of not knowing whether a proposition is true or false. Holton [Holton 2004, p. 22] provides an illustrative example which we summarize here.

Suppose you visit a casino and join an ongoing game. Someone else is about to roll a die. In this game, if this person rolls a three, then you will lose \$100. What is your risk? What is your assessment of the probability that you will lose \$100? Many might assess their chances of losing at one in six. But such an assessment carries with it the assumption of a six-sided die. In this casino game, the die is ten-sided. The salient points of this example emphasize the following: an individual is uncertain of a proposition if:

- The individual does not know whether or not the proposition is true or false.
- The individual is unaware of the proposition, itself.
- While probability is often used as a metric of uncertainty, probability can only quantify perceived uncertainty.

Now that we have defined uncertainty, we turn our attention to exposure, the other essential component of risk. Exposure relates to the material significance of the proposition. An entity (e.g., individual, department, institution) is exposed to a proposition if and only if that entity “*would care* whether or not the proposition is true.” This definition of exposure is consistent with the above discussion of perceived uncertainty, in that it emphasizes the possibility of exposure while unaware of the existence of the proposition (e.g., unknown unknowns). Holton also provides the following simple yet illustrative example of exposure: Suppose it is raining and you find yourself outside without a rain coat. You have exposure because you have a preference as to

the true state of the proposition, it is currently raining. In summary, exposure relates to those propositions which would have material consequences.

With the two essential components of uncertainty and exposure, we can now present Holton’s definition of risk as: “the exposure to a proposition of which one is uncertain.” One final clarifying example involves a man jumping from an airplane without a parachute. If the man faces certain death, then he therefore faces no risk. Risk requires both exposure and uncertainty.

2.2 Overview of Risk Metrics

A risk metric is an attempt to measure or depict aspects of perceived risk [Holton 2004]. As with a model, the purpose of a risk metric is usefulness, not verisimilitude. The famous quote from George Box [Box 1979, p. 2], past president of the American Statistical Society, seems appropriate here: “Models, of course, are never true, but fortunately it is only necessary that they be useful.” And Holton concurs with Box, that like models, it is only necessary that risk metrics be useful.

Risk metrics are values, that are produced by an evaluation method based on analysis guided by a risk model under a set of assumptions. They indicate the perceived level of risk in a manner that can be direct (the metric measures a perception of risk) or indirect (the metric measures some risk factor or set of factors, and serves as an indicator of perceived risk). Risk metrics can be produced with varying degrees of scope or aggregation, as illustrated in Figure 1. Note that the lower three degrees of aggregation correspond to the three levels of risk management implementation in the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) [NIST 2014].

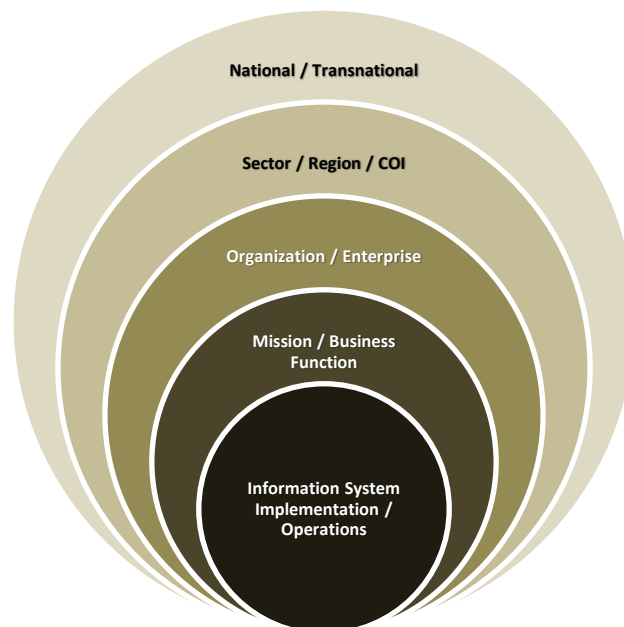


Figure 1. Risk Is Experienced at Different Scopes [Adapted from Bodeau 2014]

Risk metrics are used in multiple ways, including to aid understanding of trends (is risk increasing or decreasing?); to support decisions among alternative courses of action (also

referred to as risk responses); and to comply with legal or regulatory requirements or to address the concerns of such stakeholders as shareholders (is the level of risk acceptable?). Risk responses can be characterized in general terms (e.g., risk acceptance, risk mitigation, risk sharing, risk transfer [NIST 2011]) or in specific terms (how does this set of possible risk mitigations compare with that one?), and can be taken with different scopes.

S&T seeks cyber risk metrics to assess whether the solutions S&T provides are serving to reduce cyber risks experienced by FSS organizations, the sector as a whole, and the general public, which can be evaluated across the FSS. Ideally, all organizations in the FSS would evaluate a quantitative cyber risk metric and report it to DHS S&T, so that it could be transformed and aggregated into metrics at the sub-sector and sector levels. These aggregated metrics could be tracked; such tracking would both serve to identify trends and help determine the effectiveness of S&T-provided cybersecurity solutions.

Additionally, each organization has exposure and therefore would be motivated to evaluate its cyber risk metric. In an ideal context, that metric could also be used as part of the organization’s overall risk management and benchmarking process, easily related to the other risk metrics it uses (for example, clearly indicating how cyber risk contributes to reputation risk).

In practice, this ideal can only be approximated. Organizations differ with respect to the maturity and sophistication of their overall risk management processes and of their cybersecurity programs; these differences result in different capabilities to evaluate cyber risk metrics. Figure 2 reflects a simple depiction of these differences.

H/M/L Assessments	Semi-Quantitative Assessments	Quantitative Model-Based Metrics
Low-fidelity: <ul style="list-style-type: none"> • Based on generic, high-level model <ul style="list-style-type: none"> • Underlying assumptions not stated; only a few factors identified • Evaluated by subject matter experts <ul style="list-style-type: none"> • Low level of effort • Rapid assessment; can focus attention on areas for more detailed assessments • Typically qualitative, may be related to ranges of semi-quantitative values 	Tailorable-fidelity: <ul style="list-style-type: none"> • Driven by underlying model <ul style="list-style-type: none"> • Underlying assumptions are often implicit • Evaluated in notional or representative environment <ul style="list-style-type: none"> • Tailorable level of effort, but generally moderate rather than low • Can focus attention on factors to be evaluated with greater fidelity • Sensitive to expertise, experience of evaluators • Typically semi-quantitative, or quantitative with low granularity 	High-fidelity: <ul style="list-style-type: none"> • Rely on explicit model <ul style="list-style-type: none"> • Dependencies among factors identified • Evaluated in well-defined, calibrated environment <ul style="list-style-type: none"> • Level of effort depends on availability of tools for information-gathering and analysis – often high for first evaluation, moderate-to-low thereafter • Sensitivity to inputs, parameters, assumptions can be determined • Produce quantitative values • Can be supported by evidence (e.g., indicators, observables)

Figure 2. Spectrum of Risk Metric Approaches

Organizations differ in their risk framing [NIST 2011]: which threats they believe they face, as well as which types of risk are of greatest concern; these differences cause them to measure different risk factors. Organizations differ in exposure and risk tolerance; since those with higher risk tolerance may need less fidelity in risk measurement than those with lower risk tolerance,

some organizations may be content with qualitative assessments while others seek quantitative values.

2.3 Risk Metrics in the Financial Services Sector (FSS)

In this section, we briefly describe risk metrics in the FSS from the perspectives of three distinct classes of stakeholders: government, industry and academia.

2.3.1 Government

Government efforts to develop and track FSS risk metrics are led by organizations such as the Office of Financial Research (OFR) and the Federal Financial Institutions Examination Council (FFIEC). Below, we provide a brief overview of these organizations and their respective efforts relating to FSS risk metrics.

2.3.1.1 Office of Financial Research (OFR)

Since its establishment in 2010 under the Dodd-Frank Wall Street Reform and Consumer Protection Act, the OFR collects FSS data and conducts FSS related research at the direction of the Financial Stability Oversight Council (FSOC). The entire purpose of the OFR is “to promote financial stability by looking across the financial system to measure and analyze risks, perform essential research, and collect and standardize financial data.”¹ Therefore, the OFR exists to develop and apply risk metrics at the FSS level. These metrics are depicted throughout OFR tools such as the [Financial Markets Monitor](#), the [Financial Stability Monitor](#) as well as in the [OFR Financial Stability Reports](#). It should be noted that while the OFR has undertaken extensive efforts to develop and apply risk metrics aimed at the stability of the U.S. financial system, none of the above tools and analyses reflect cybersecurity risks. In its most recent Financial Stability Report, OFR [OFR 2015] points to the potential usefulness of the FFIEC’s Cybersecurity Assessment Tool, which we describe below.

2.3.1.2 Federal Financial Institutions Examination Council (FFIEC)

Dating back to its establishment in 1979, the FFIEC² is a formal interagency body comprised of:

- Board of Governors of the Federal Reserve System (FRB)
- Federal Deposit Insurance Corporation (FDIC)
- National Credit Union Administration (NCUA)
- Office of the Comptroller of the Currency (OCC)
- Consumer Financial Protection Bureau (CFPB)

FFIEC performs two primary functions:

¹ Office of Financial Research website: <https://financialresearch.gov/about/>

² Federal Financial Institutions Examination Council website: <https://www.ffiec.gov/about.htm>

- prescribes uniform principles, standards, and report forms for the federal examination of financial institutions
- makes recommendations to promote uniformity in the supervision of financial institutions

The FFIEC IT Examiners Handbook defines risk and risk measurement as follows:

Risk: The potential that events, expected or unanticipated, may have an adverse effect on a financial institution's earnings, capital, or reputation.

Risk measurement: A process to determine the likelihood of an adverse event or threat occurring and the potential impact of such an event on the institution. The result of risk measurement leads to the prioritization of potential risks based on severity and likelihood of occurrence.

Thus, measurement of risk has two major aspects: measurement of the likelihood of occurrence of an adverse event (which in turn is a function of the properties of the threat and systemic properties such as vulnerability) and measurement of the severity of the consequences or impact.

2.3.2 Industry

For brevity and relevance, we describe a broad category of probabilistic risk measures, known as Value at Risk (VaR). Prior to the financial crises of 2008, VaR metrics were widely considered to be the prevailing risk metrics in financial services industry. In this section, we define VaR and its applications; and below, we describe the theoretical origins of VaR.

2.3.2.1 Definition of Value at Risk (VaR)

A financial portfolio consists of several, distinct assets. Each asset has a current market value; therefore, the portfolio also has a current market value. At some point in the future (e.g., day, week, month, year), the market value of the assets is uncertain, and therefore the value of the portfolio is also uncertain. We can designate a random variable to represent the uncertainty in the future value of the portfolio, and we can assign a probability distribution to that random variable. A VaR metric³ is a function of both the current market value of the portfolio and the probability distribution [Holton 2014]. Given a portfolio and any VaR metric, a VaR measure would be the application of any function which determines specific values within the portfolio for that metric.

2.3.2.2 Applications of Value at Risk (VaR)

While VaR could be germane to any application, including exposure to market risk, we provide an overview of some common uses of VaR in the financial services industry. It can be used in a variety of different contexts including the evaluation of risk associated with trading and lending activities carried out by banks and the evaluation of risk in investment portfolios [Beckstrom 2014, Buith 2015]. VaR is also used in "computing capital adequacy requirements" for banking

³ Here we adopt NIST's distinction between *metrics* and *measures*, as put forth by Black [Black 2008]. A metric is a tool used to facilitate decision making, and to improve performance and accountability. A measure consists of quantifiable, observable, and objective data, and measures are used to support metrics.

institutions around the world [Beckstrom 2014]. The metric produced by VaR can be interpreted as follows:

“For a given time horizon t and confidence level p , the value at risk is the loss in market value over the time horizon t that is exceeded with probability $1-p$.” [Duffie 1997]

2.3.2.3 Risk Limits

Throughout the financial services industry, any financial institution faces the two essential components of risk: exposure and uncertainty. Therefore, these organizations might attempt to determine what risks are acceptable and what are not. These bounds around risk are known as risk limits, as well as risk appetite. A risk limit has three components: (1) a risk metric, (2) an associated risk measure, and (3) a bound, which is the value for the risk metric that is not to be breached. An organization might institute a mandatory reporting requirement whenever a risk limit is exceeded.

2.3.2.4 Risk Reporting and Oversight

Closely related to risk limits, above, risk reporting and oversight reflect a policy or formal program requiring that certain risk limits be followed and violations of those limits be reported. While such a policy may be instituted within an institution, formal risk reporting and oversight programs are often administered by regulatory agencies. Below, we describe an example of such a program, namely, capital requirements.

2.3.2.5 Capital Requirements

Bank regulatory capital requirements provide an illustrative example of mandatory risk reporting and oversight. Rather than regulators specifying which banking activities are deemed acceptable and which unacceptable, regulators instead opted to require that banks hold capital reserves to cover potential losses resulting from their risk-taking activities. This set of bank regulatory capital requirements became known as the “Basel Accord”. In 1996, BCBS released an amendment to the Basel Accord “to apply capital charges to the market risks incurred by banks” [BCBS 1996a]. These capital charges were to be computed “according to one of two methods, a standardised measurement method or an approach based on the results of internal models” [BCBS 1996a]. Among the quantitative criteria governing the method based on internal models is a requirement “that ‘value-at-risk’ be computed daily, using a 99th percentile, one-tailed confidence interval” [BCBS 1996b].

2.3.3 Academia

Here, we provide a brief overview of the theoretical foundation of VaR which follows along two parallel theoretical trajectories: portfolio theory and calculations of capital adequacy.

2.3.3.1 Portfolio Theory

Widely credited to Markowitz [Markowitz 1952], portfolio theory describes an approach to compiling a portfolio of assets such that the expected return of the portfolio is maximized for a given level of acceptable risk. The key contribution of portfolio theory is that the risk and return of a specific asset should not be assessed at the level of the particular asset, but by how the asset

contributes to the overall risk and return of the portfolio. Markowitz would later earn a share of the 1990 Nobel Prize in Economics for his 1952 paper on portfolio theory.

2.3.3.2 Calculations of Capital Adequacy

Before the regulations enacted in 1933, the U.S. securities markets were self-regulated. As early as 1922, the New York Stock Exchange imposed capital requirements on its members [Dale 1996]. Following the 1929 stock market crash, the U.S. implemented various legislation to restore investor confidence in the financial system.

The 1933 Banking Act, also known as the Glass – Steagall Act, established federal deposit insurance in addition to a segregation of the banking and securities industries. The 1933 Banking Act made the clear distinction between commercial banking (i.e., the business of taking deposits and making loans) and investment banking (i.e., the business of underwriting and dealing in securities). And the 1934 Securities Exchange Act established the Securities Exchange Commission which immediately gained regulatory oversight to set capital requirements on investment banks and other securities firms.

From the perspective of internal programs for capital adequacy, Garbade [Garbade 1986, 1987], who is now a senior researcher at the Federal Reserve Bank of New York, is often credited for the sophistication of his early VaR measures for assessing internal, capital requirements. Garbade attempted to develop an approach to risk assessment and capital adequacy which was simultaneously comprehensive yet simple. Instead, he concluded that “risk and capital adequacy formulas are either complex or of limited applicability, and are sometimes both.”

In Garbade’s 1986 paper, he developed VaR measures which depicted each asset in a portfolio based upon the asset’s price sensitivity to changes in yield, assuming that future portfolio market values were Gaussian distributed. Garbade [Garbade 1987] developed a technique to disaggregate a portfolio’s risk and distribute it across multiple, profit centers. Note that Garbade’s attempts to assess risk and to determine capital adequacy preceded those of the Basel Accords by a decade or more.

3 Cybersecurity Risk Metrics

The goal of this paper is to develop a way to roll up various idiosyncratic approaches at the institutional level to a sector level generalizable view for DHS S&T to better identify and assess the impact of cybersecurity related investments. While the initial focus is on the FSS because of its level of sophistication in comparison to other sectors – see Section 4 for details – future iterations will attempt to extend the findings to other sectors.

As discussed below, there are a variety of cyber risk metrics in use, which have in turn been developed using a wide range of approaches at various levels of maturity. The state of the current practice includes: risk metric development frameworks, quantitative approaches, and practice-oriented approaches such as interactive/iterative red teaming exercises.

3.1 Risk Metric Development Frameworks

Frameworks for developing risk metrics include risk models such as the general model offered by NIST SP 800-30R1 [NIST 2012] and more targeted frameworks, focused on a specific aspect of risk, such as the cyber investment framework which focuses on vulnerabilities. Frameworks are often designed to be extensible or tailorable to reflect the orientation of varying organizations. Therefore, framework-based metrics are often qualitative or semi-quantitative, to enable organizations to provide anchoring examples for the values they provide and to enable additional factors to be defined and combined easily.

3.1.1 Risk Models

A risk model typically provides a language for talking about risk, i.e., about the potential for adverse consequences. Risk is usually modeled in terms of the type, likelihood, and severity of those consequences, using some set of risk factors. As discussed in more detail in the Threat Model Survey, NIST SP 800-30R1 offers a general model of cyber risk. It identifies as key factors threat sources, threat events (or threat scenarios, built out of a set of threat events), vulnerabilities and predisposing conditions (systemic properties which increase the likelihood that a threat event or threat scenario will result in adverse consequences), likelihood of occurrence, likelihood of impact, consequences, and impact severity.

NIST SP 800-30R1 is designed to enable each organization to define the risk models and corresponding risk metrics that best fit into its overall enterprise risk management process. For illustrative purposes, NIST SP 800-30R1 does define value scales for various aspects of risk. These value scales provide qualitative values (very low to very high) as well as semi-quantitative scales (0-10 or 0-100).

The Object Management Group (OMG) has issued a Request for Proposal for an operational Threat and Risk Model [OMG 2014]. The requested model is expected to include concepts related to operational threats and risks, as well as additional granularity and detail in the cyber

domain. Portions of the requested model are expected to be mapped to Structured Threat Information eXpression (STIX)⁴.

3.1.2 Cybersecurity Investment Framework

Despite the number of cybersecurity and risk assessment models developed over the years, because of the lack of credible data about the number and impact of cyber attacks, it is very hard for organizations to quantify the risks they face. There is also almost uniform agreement that there is no generally accepted model or set of investment principles that organizations can use to guide cybersecurity investments. As a result of the lack of solid, usable cyber risk metrics, organizations fall back on metrics characterizing their vulnerabilities and their security controls, management, monitoring, and response capabilities. These in turn are used to identify gaps requiring investment and set the basis by which those investments are evaluated. In an effort to help develop a useful investment framework, MITRE participated in a working group organized by the Armed Forces Communications and Electronics Association (AFCEA)⁵ which published two papers [AFCEA 2013, 2014] on the Economics of Cybersecurity. The papers recommended both a framework and principles for making and assessing investments in cybersecurity, along with a set of proposed quantifiable metrics.

The recommendations in both these papers were reconfirmed in a series of interviews with chief information security officers (CISOs) and senior cybersecurity risk management professionals from across the FSS. The figures below represent the key findings in the AFCEA papers.

The AFCEA 2013 paper introduced a practical framework (Figure 3) for guiding cybersecurity investments based in part on the observation that the majority of damaging cyber attacks were fairly unsophisticated and documented evidence that a baseline of security controls can be effective against most of these types of attacks. There were a number of recommendations around the implementation of the framework, which included the adoption of a set of critical controls. The graphical model presented in Figure 3 is fairly easy to understand and helped articulate several important principles that should guide cybersecurity investments. These include:

- Implementation of a comprehensive baseline of security controls that address threats that are of low to moderate sophistication is essential and is economically beneficial.
- Focus security investment beyond the baseline controls to counter more sophisticated attacks against the functions and data that are most critical to an organization.
- For sophisticated attacks, an organization should accept the security risk of not protecting functions and data that are of lowest impact to the organization's mission and where cost exceeds benefits.

⁴ STIX (Structured Threat Information eXpression) is a standardized XML programming language for conveying data about cybersecurity threats in a common language that can be easily understood by humans and security technologies.

⁵ The Armed Forces Communications and Electronics Association (AFCEA) is a non-profit membership association that was established in 1946 as a forum for advancing professional knowledge and relationships in the fields of communications, information technology, intelligence, and security.

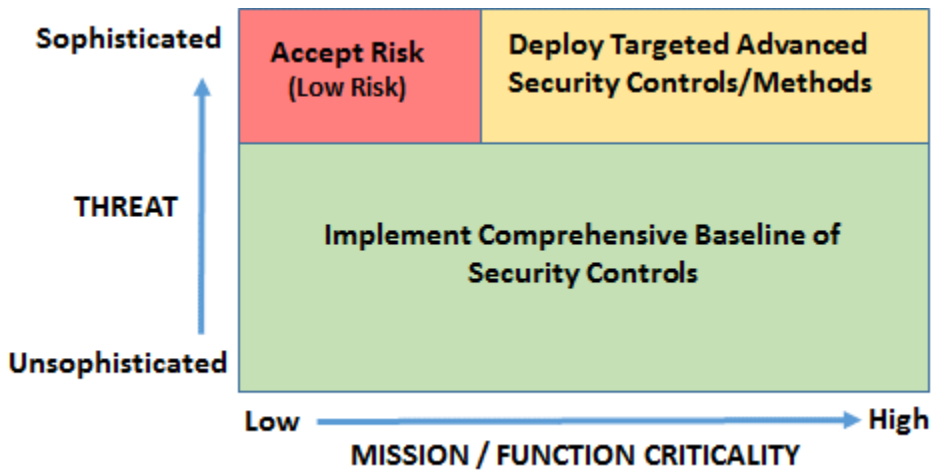


Figure 3. Cybersecurity Economic Framework [Adapted from AFCEA 2013]

The AFCEA 2014 paper extended this basic Cybersecurity Economic Framework to provide guidance to organizations on how to address more sophisticated cyber threats⁶. In developing the Extended Cybersecurity Framework depicted in Figure 4, a number of cybersecurity models and methodologies were examined. The extended framework reflects a maturity model that includes the elements of the initial framework in the right column. The extended graphical model presented in Figure 4 in turn helps articulate several more important principles that should help organizations appropriately focus and prioritize investments. These include:

- The economic benefit of participating in multiple, high quality cybersecurity information sharing exchanges regarding the dynamic characteristics of sophisticated threats is very high.
- Additional investments to address sophisticated threats should be specifically tailored to the (evolving) threat characteristics.
- Effective countering of the most sophisticated threats (e.g., nation-state) requires investment in current technology controls and human capabilities to be able to effectively predict and respond to attack patterns.

⁶ A threat to cyber systems refers to persons who attempt unauthorized access to a cyber system device and/or network. This access can be directed either from within an organization by trusted users, or from remote locations by unknown persons using the Internet. Threats can come from numerous sources, including hostile governments, terrorist groups, disgruntled or poorly trained employees, and malicious intruders.

<u>Enhanced Descriptor</u>	<u>Employment of Security Controls</u>	<u>Security Tailored to Mission</u>	<u>Participate in Information Sharing (threat & vulnerabilities)</u>	<u>Response to Cyber Threats</u>	<u>Cybersecurity Framework Area</u>
Level 4: Resilient Operate through Sophisticated Attack	Augment CSC Based on Mission & Threats	Investments are Mission Assurance focused	Tools and Staff to Respond to Shared Threat Information	Analytical Capabilities to Anticipate Threats	Additional Investments to Deploy Targeted Advanced Security Controls/Methods
Level 3: Dynamic Able to respond to Sophisticated Attack	Augment CSC Based on Mission & Threats	Investments are Mission Protection focused	Tools and Staff to Respond to Shared Threat Information	Capabilities for Rapid Reaction to Threats	
Level 2: Managed Protection against Unsophisticated Attack	CSC Integrated and Continuously Monitored	Partially Mission focused	Respond to Information Inputs	Respond to Attacks after the Fact	Implement Comprehensive Baseline of Security "Good Hygiene"
Level 1: Performed Some Protection Against Unsophisticated Attacks	Foundational/ Critical Security Controls (CSC) Implemented	Mission Agnostic	Inconsistent Response to Information Inputs	Respond to Attacks after the Fact	

Figure 4. Extended Cybersecurity Framework [Adapted from AFCEA 2014]

Finally, during the development of the frameworks, several participants offered examples of a number of quantifiable metrics that could be used for determining the effectiveness and economic benefits of cybersecurity measures. While the paper did not focus on the development of specific metrics, it was determined that the framework did lend itself to an initial set of metrics with the recommendation that much more work needs to be done [AFCEA 2014]. Figure 5 extends the framework to include this initial set of metrics.

Note that this initial set of metrics measures properties of an organization’s cybersecurity program, enterprise architecture, and operational processes and procedures. Higher values reflect a greater ability to address cyber risks. Lower values reflect systemic weaknesses or vulnerabilities.

Enhanced Descriptor	Employment of Security Controls	Security Tailored to Mission	Participate in Information Sharing (threat & vulnerabilities)	Response to Cyber Threats
Level 4: Resilient Operate through Sophisticated Attack	Metric: Capability for real time deployment of controls in response to changing threat profile	Metric: 1) Deployed protection architecture based on assuring mission continuity; 2) Regular exercise of ability to operate through attack	Metric: 1) Robust network of information exchange partners monitored on real time basis; 2) Staff capable of extending threat data to predict threat evolution	Metric: Established policies and practices as well as experienced staff able to permit real time responses to sophisticated threats
Level 3: Dynamic Able to respond to Sophisticated Attack	Metric: Implement threat monitoring capabilities to support identification and deployment of additional controls	Metric: 1) Identification of mission critical capacities; 2) Deployment of partial architecture and controls to protect mission critical capabilities	Metric: 1) Robust network with information exchange sources; 2) Experienced staff capable of rapid response to sophisticated threats	Metric: Organic staff capable of recognizing sophisticated threats and recommending response actions
Level 2: Managed Protection against Unsophisticated Attack	Metric: 1) Ensure baseline controls are consistently applied across the enterprise; 2) controls are implemented with continuous automated monitoring with a goal of hourly or single digit minute cycle times	Metric: Formal identification of mission critical capabilities	Metric: 1) Established relationship with one or more information sources for cyber threat and vulnerability information; 2) Standard processes for rapidly responding to threat/vulnerability updates	Metric: Organization staff able to respond after the fact to attack
Level 1: Performed Some Protection Against Unsophisticated Attacks	Metric: 1) Implement DND top 4 Controls; 2) Implement some additional CSC or DND 35 Controls	Metric: None	Metric: Threat / Vulnerability information pushed to organization but inconsistently reviewed or applied	Metric: Attack response prompted from outside the organization

Figure 5. Metrics for Enhanced Cybersecurity Framework [Adapted from AFCEA 2014]

3.1.3 Framework-Based Tool

The FFIEC developed the Cybersecurity Assessment Tool (CAT) to help institutions identify their risks and determine their cybersecurity preparedness. The purpose of the CAT is to provide “a repeatable and measurable process for financial institutions to measure their cybersecurity preparedness over time” [FFIEC]. The CAT incorporates FFIEC guidance on cybersecurity with the NIST CSF.

3.2 Quantitative Approaches

As observed above, frameworks often provide a foundation for defining qualitative or semi-quantitative risk metrics, based on extensible or tailorable risk models. However, more detailed risk models can be defined with the goal of supporting quantification and computation. Two which are of particular relevance to the financial sector are Cyber VaR (CVaR) and Factor Analysis of Information Risk (FAIR).

3.2.1 CVaR

In recent years, efforts have been made to modify VaR in order to create methods specifically designed to quantify cybersecurity risk. These new methods, referred to as Cyber VaR, provide “top management with a single risk number and a statistical probability to understand the overall cyber security risk of an enterprise” [Beckstrom 2014]. Cyber VaR has two primary goals. The first is to “help risk and [information security] professionals articulate cyber risk in financial

terms” [FAIR]. The second is to “enable business executives to make cost-effective decisions and achieve a balance between protecting the organization and running the business” [FAIR]. Although Cyber VaR methods show strong potential in their ability to quantify cybersecurity risk and streamline investment decisions to counter that risk, they also have a number of limitations. Due to the fact that Cyber VaR methods constitute a relatively new form of risk evaluation, at least some methods lack the refinement and field testing necessary for deployment [Buith 2015].

Further limitations of Cyber VaR include its reliance on historical data which is not always available and not necessarily predictive of the future, as well as the inherent difficulty faced when attempting to quantify all vulnerabilities an adversary could attack [Buith 2015]. Additionally, Cyber VaR’s effectiveness is limited by a lack of standard maturity frameworks and the fact that Cyber VaR supports only a limited number of risk scenarios⁷. Despite these limitations, Cyber VaR methods show great promise in their ability to impose a structured and quantitative approach to risk management in an area that has historically lacked such an approach [Reagan 2016], as such it is also a key element of the World Economic Forum’s approach to improving cyber resilience [WEF 2015].

3.2.2 FAIR

Factor Analysis of Information Risk (FAIR) is a Cyber VaR method chosen as the “international standard information risk management model” by the Open Group, a global consortium with over 500 member organizations including HP, IBM, Oracle, and MITRE [FAIR]. FAIR is defined as “a standard Value-at-Risk model for information and operational risk that helps information risk, cybersecurity and business executives measure, manage and communicate on information risk in a language that business understands, dollars and cents” [FAIR]. In order to communicate the principles underlying FAIR, the Open Group produced the Open FAIR Body of Knowledge, which consists of two standards [Josey 2014]. The first standard, known as the Open Risk Taxonomy, “defines a taxonomy for the factors that drive information security risk” [Josey 2014]. The second standard, Open Risk Analysis, “describes process aspects associated with performing effective risk analysis” [Josey 2014]. Together, these standards allow organizations to “Speak in one language concerning their risk”, “Consistently study and apply risk analysis principles to any object or asset”, “View organizational risk in total”, and “Challenge and defend risk decisions” [Josey 2014].

Within the Open Risk Taxonomy, risk is broken into two primary components [Josey 2014]. The first component, loss event frequency, is further divided into threat event frequency and vulnerability factors. Threat event frequency and vulnerability are then divided into additional factors. The second component, loss magnitude, is further divided into primary loss and secondary loss factors. Secondary loss is then divided into additional factors.

Open Risk Analysis utilizes a number of established quantitative techniques including “PERT distributions, calibrated estimates to support data, and Monte Carlo stochastic simulation” [Josey 2014]. In order to quantify cybersecurity risk in a manner consistent with FAIR, a company

⁷ An IT risk scenario is a description of an IT related event that can lead to a business impact, when and if it should occur.

known as RiskLens has produced a number of software applications [RiskLens]. Either these or other risk management applications would be necessary to implement FAIR in an organization.

Overall, FAIR appears to be a logical and structured Cyber VaR method that is mature enough for deployment within an organization. As a deployment of FAIR would likely require the acquisition of software from a vendor, it is imperative that the specific software applications be vetted on a technical level to ensure that the quantitative techniques and assumptions align with best practices from academic literature and industry.

3.2.3 Industry Developed Metrics

There have been a number of metrics developed by individual institutions and consulting organizations over the years. Based on the interviews discussed in more detail in Section 4, they all basically attempt to measure risk as:

$$\text{Risk} = \text{Vulnerability} \times \text{Threat} \times \text{Asset Value} \times \text{Probability of Occurrence}$$

Almost all the interviewees acknowledged that it is extremely difficult, if not impossible, with current limited sources of data and the general immaturity of cyber risk assessment approaches, to actually assign a quantitative or dollar amount to any of the items in the formula. They look to the elements to help guide their research and inform their decisions, which includes investments in resilience initiatives to mitigate the impact or consequence of an attack. The current focus is on reducing the “probability of occurrence” as the primary way to assess the value of investments, which is in turn covered in the threat modeling processes discussed in the accompanying *Survey of Threat Model* document. To that end, information security professionals follow similar processes as that published by SANS in Payne [Payne 2006]:

1. Define the metrics program goal(s) and objectives
2. Decide which metrics to generate
3. Develop strategies for generating the metrics
4. Establish benchmarks and targets
5. Determine how the metrics will be reported
6. Create an action plan and act on it, and
7. Establish a formal program review/refinement cycle

They have all also implemented either a Top-Down or Bottom-Up approach similar to that described in Payne [Payne 2006].

Finally, lacking cyber risk metrics, institutions rely on metrics of security compliance, incidence, and response that include the following areas: process, network, software, and people related security. An example is the CIS Security Metrics developed by the Center for Internet Security [CIS 2009], which includes consideration of the function that is being assessed, what management is interested in knowing or assessing, and a set of related sample metrics. With the exception of Incident Management, these metrics focus on vulnerability, since reducing vulnerability is a well-established strategy for reducing risk. However, reducing consequences is also an important strategy, one which requires an assessment of consequences. Asset value (used

in the formula at the start of this subsection) is a key input into assessment of consequence severity, but is not necessarily identical to that value. Asset value assessment is an ongoing topic of investigation [Miller 2016].

3.3 Red Teaming

In order to evaluate the effectiveness of their cyber defenses, many organizations engage in red teaming activities. The Information Design Assurance Red Team (IDART) at Sandia National Laboratories defines red teaming as “authorized, adversary-based assessment for defensive purposes” [Sandia]. One of the key differences between red teaming and other techniques for evaluating cyber defenses is that red teaming focuses on what “one or more adversaries [would] do if they were attacking [a] target” and considers an adversary’s “knowledge, skills, commitment, resources, and culture” [Sandia]. Red teaming can be used as a tool to assist organizations in carrying out a number of functions including “Anticipate program risk”, “Support security decisions”, “Identify and describe consequential security design alternatives”, and “Measure security progress and establish security baselines” [Sandia]. Data collected through red teaming activities can be used to inform metrics that aid in the understanding of an organization’s cybersecurity risk.

4 Cybersecurity Risk Metrics Survey of FSS

Given the dynamic state of information security vulnerabilities and the volume and complexity of existing regulatory and compliance requirements, financial institutions face a huge challenge in modeling related threats and risks. To help understand these challenges, cybersecurity executives from a cross section of the largest financial institutions and market utilities were interviewed. Based on the results of the survey, the key questions that all the institutions face are: what is the most likely threat against specific information or business assets, and what would the dollar impact be to the firm if the threat occurred?

As described in more detail below, all the institutions have defined and documented processes, developed over time based on the experience of their CISOs and their understanding of the inherent risks of their various business functions and assets. All of them state they have well-funded operations with very experienced (albeit hard to find and retain) staffs dedicated to developing their risk assessments. However, at the same time they almost uniformly recognize their dependence on relatively immature, qualitative approaches and the need to find a way to share the data necessary to inform the development of more quantitative metrics and mature risk management capabilities.

4.1 Process

All the institutions developed processes and frameworks that included strategic planning and reviews, developing risk metrics and measures, and modeling threats. However, none of the institutions had adopted any of the published risk management frameworks. They each have adopted hybrids based on their experience, the culture of their firms, and various inputs from BCBS 329, COBIT, FAIR, FFIEC/OCC, International Organization for Standardization (ISO), Lockheed Martin Cyber Kill Chain, NIST, OCTAVE, SANS, STRIDE, and SOX. In addition, consulting firms are retained to develop customized approaches.

In general, the approach to risk modeling incorporates primarily subjective assessments of threats and vulnerabilities, with some attempts at quantification around consequence (e.g., cost of breach). The processes start with attempts to baseline current environments, look at current and evolving threats, align with future goals, identify control or capability gaps, prioritize and plan related investments, monitor execution, and incorporate a feedback or refresh process. Threats are modeled as to the objective or impact being sought using a list of actors that includes which actors are interested in various objectives, and whether or not the actor has the capabilities to achieve the desired impact.

At a strategic and board of directors level the process is to manage to a predefined, but qualitatively based risk appetite that is periodically revised. Investments are primarily focused where controls are deemed inadequate to achieve or maintain the desired risk levels. Root cause analysis is used to determine the source of failures, and potential gaps in controls, along with external audits and pen testing to assess the adequacy of the overall risk management process. There are no traditional return on investment or financial calculations, and budgets are primarily built from the ground up based on the CISO's recommendations. For example, it is very difficult to answer a question about how much a firewall is worth; however very few would argue that an organization should not have one in place.

4.2 Data

The focus of their efforts is primarily around the adequacy of the controls that are in place – both protective and detective – and any incident data that may be related to them. Besides incidents, most organizations maintain risk, threat, and control libraries which are updated as needed from either experience or a variety of third party sources. Some of the third-party sources referenced included Bitsight, Whitehat, Gartner, Financial Services Information Sharing and Analysis Center (FS-ISAC), iSight Partners, and CERT, which provide a variety of alerts, compromises, threats, and end point and network monitoring data. Beyond counts related measures, there are really no quantitative metrics.

Almost all the interviews indicated there is no meaningful, standardized data that can be collected and ingested on an automated basis, or that adequately covers developing risks due to new capabilities or motivations. There are some operating risks that have meaning, such as loss avoidance type measures (e.g., frauds identified and stopped), but nothing they thought could be used for traditional return on investment (ROI) or cyber defense performance metrics. For the most part, calculations of risk appetite and losses come down to professional judgment and will remain there until the necessary data is available. In addition, most pointed to the inadequacy of their user behavioral, network behavioral, and baselining data capture and analysis capabilities.

4.3 Reporting

For the most part, the information security budget is an aggregation of the costs related to control and other recommendations. As discussed above, there is really no quantitative evaluation that is considered effective or adequate in either setting the budget or in assessing the return on related investments. There is some benchmarking against peer groups around percent of total revenues and percent of total full-time equivalent, or FTE (both of which are roughly 1%), but the investment and budget is primarily driven by the credibility of the CISO. It was mentioned that the regulatory community uses “security dollar spent per headcount” as a metric and then asks institutions to justify if they spend less than their peers.

Almost all the organizations struggle with how best to communicate with their board, and generally steer away from information briefs and focus on the justification for specific decisions the CISO would like them to make. Because it is hard to tie specific outcomes to specific investments, the focus is on items such as trends in team response effectiveness and related improvements in efficiency, as well as heat maps and variances from predefined risk levels based on a variety of qualitative assessments. Most also incorporate some form of control and capability maturity assessments and include trends related to their growth in their reports.

Threat level data is aggregated to a macro level, is updated monthly, and includes considerations around culture, compliance, and incident response. They consider incident likelihood, incident locations, present mitigation status, and major initiatives with a heat map. This includes top 10 type lists related to items such as technology life cycle, architecture, malware, regulation, capability, and control framework. Board level discussions center around examining performance versus forecasts, increasing efficiencies at lower investments, and benchmarking against peers.

4.4 Testing & Validation

Most of the firms use external firms to conduct some form of red teaming, as well as penetration testing and compliance audits. Sometimes the business units are involved, but for the most part the businesses are responsible for the use of their applications and the policies they create to mitigate related risks while the information security group is responsible for IT controls and their effectiveness. In addition, all the firms participated in or conducted war gaming exercises of one form or another.

The internal exercises are primarily focused on making sure cybersecurity is in alignment with the business units, preventing destructive malware aimed at customer data, mitigating the impact of distributed denial of service (DDOS) attacks, and minimizing financial theft. The list of external exercises included Cyber Guard, Hamilton Alliance, Quantum Dawn, a variety of regional and global exercises, and one-off exercises conducted between peer organizations. For those that choose not to participate in an external event, they can usually get the results for those that include the FS-ISAC.

During the interviews, almost all the firms made comments on the need to improve some aspect of the existing war gaming exercises. Suggestions included creating industry only exercises, extending the participants both globally and across sectors, focusing more on scenarios related to external dependencies (e.g., Society for Worldwide Interbank Financial Telecommunication [SWIFT] and market utilities) and resiliency, and consolidating and standardizing delivery. General concerns were expressed about having regulators participating due to both their unwillingness to share how they would respond to the events in the exercise, and their tendency to include what they learn in subsequent examinations.

5 Synthesis of Cybersecurity Risk Metrics Findings

While a subset of the metrics identified during the executive interviews are assessed in more detail below, our preliminary observations include:

- In general, the various external frameworks are not considered very good, either due to an incomplete taxonomy, the complexity and cost related to understanding and implementing, their reliance on a qualitative versus quantitative methodology, and the lack of meaningful automation or external data to validate.
- While the lack of meaningful, quantitative data was a constant theme, as was the desire to have better visibility into their environments and anomalous behaviors, it was unclear that there has been any real internal or external analysis of the data necessary to improve the accuracy or completeness of the risk management practices. As a result, the models are based on subjective probability assessments, along with an acknowledged decision to violate basic math by aggregating and trending ordinal scales.
- In addition, it is unclear whether the predictions that are made related to the likelihood or impact of events are analyzed to assess the value of the current approaches. However, there was general agreement that if the right mechanism could be created, and their legal advisors would approve, they would share this data for benchmarking and analysis purposes, including the assessment of the perceived benefit of, or confidence in, existing risk management practices.
- Finally, the following Analysis and Assessment section explores what metrics, if any, from existing institutional level approaches can be extrapolated to a sector-wide approach sufficient to help DHS S&T better identify, prioritize, and evaluate the impact of its research and development investments.

6 Analysis and Assessment

Interviews with executives from a cross section of the largest financial institutions and market utilities identified a number of cybersecurity metrics that are used to inform decision making at the organization level. In order to evaluate the ability of these metrics to benefit their originating organizations and the NGCI program, an assessment methodology was developed.

6.1 Assessment Methodology

Once the executive interviews were conducted, the first step in the assessment methodology was to identify cybersecurity metrics being used at the surveyed financial institutions and market utilities. When this initial set of cybersecurity metrics was identified, the metrics were then grouped into four broad families based on their characteristics. Within each family, metrics viewed as useful or representative of metrics within the family were selected for assessment. These selected metrics were then evaluated against three criteria intended to determine the benefit provided by these metrics to their originating organizations and the potential benefit these metrics could provide to the NGCI program. Figure 6 provides a visual representation of the Cybersecurity Metric Assessment Methodology.

6.2 Families of Cybersecurity Metrics

In order to organize the large number of metrics identified during the executive interviews, four cybersecurity metric families were created. These families were defined to be mutually exclusive and collectively exhaustive. Therefore, any given metric identified during the interviews would belong to exactly one family.

The first cybersecurity metric family, Quantitative Risk Metrics, consists of metrics that provide numeric representations for the risk posed by cyber-attacks to a financial institution or market utility. These numeric representations can take the form of dollar amounts or other units. Quantitative risk metrics are often produced using models that break risk into its contributing factors. Generally speaking, a reliable quantitative risk metric would be considered one of the most useful tools for cybersecurity decision making. However, the accuracy of such a metric is highly dependent on the availability and quality of input data sources.

The second family, known as Quantitative Non-Risk Metrics, includes numeric metrics that help a financial institution or market utility maintain cybersecurity situational awareness without directly estimating risk. These metrics can take the form of dollar amounts or other units. They can often be developed or tailored to address a specific cybersecurity concern of the originating organization.

The third family, known as Qualitative Risk Metrics, includes non-numeric metrics that help an organization understand and depict its level of cybersecurity risk. Often, these metrics can utilize categorical representations of risk such as low, medium, high, or critical. Qualitative risk metrics can be particularly useful when attempting to communicate risk and related trends at a high level to a corporate board.

The final cybersecurity metric family, Qualitative Non-Risk Metrics, includes non-numeric metrics that enable a financial institution or market utility to maintain situational awareness of its cybersecurity environment without directly estimating risk.

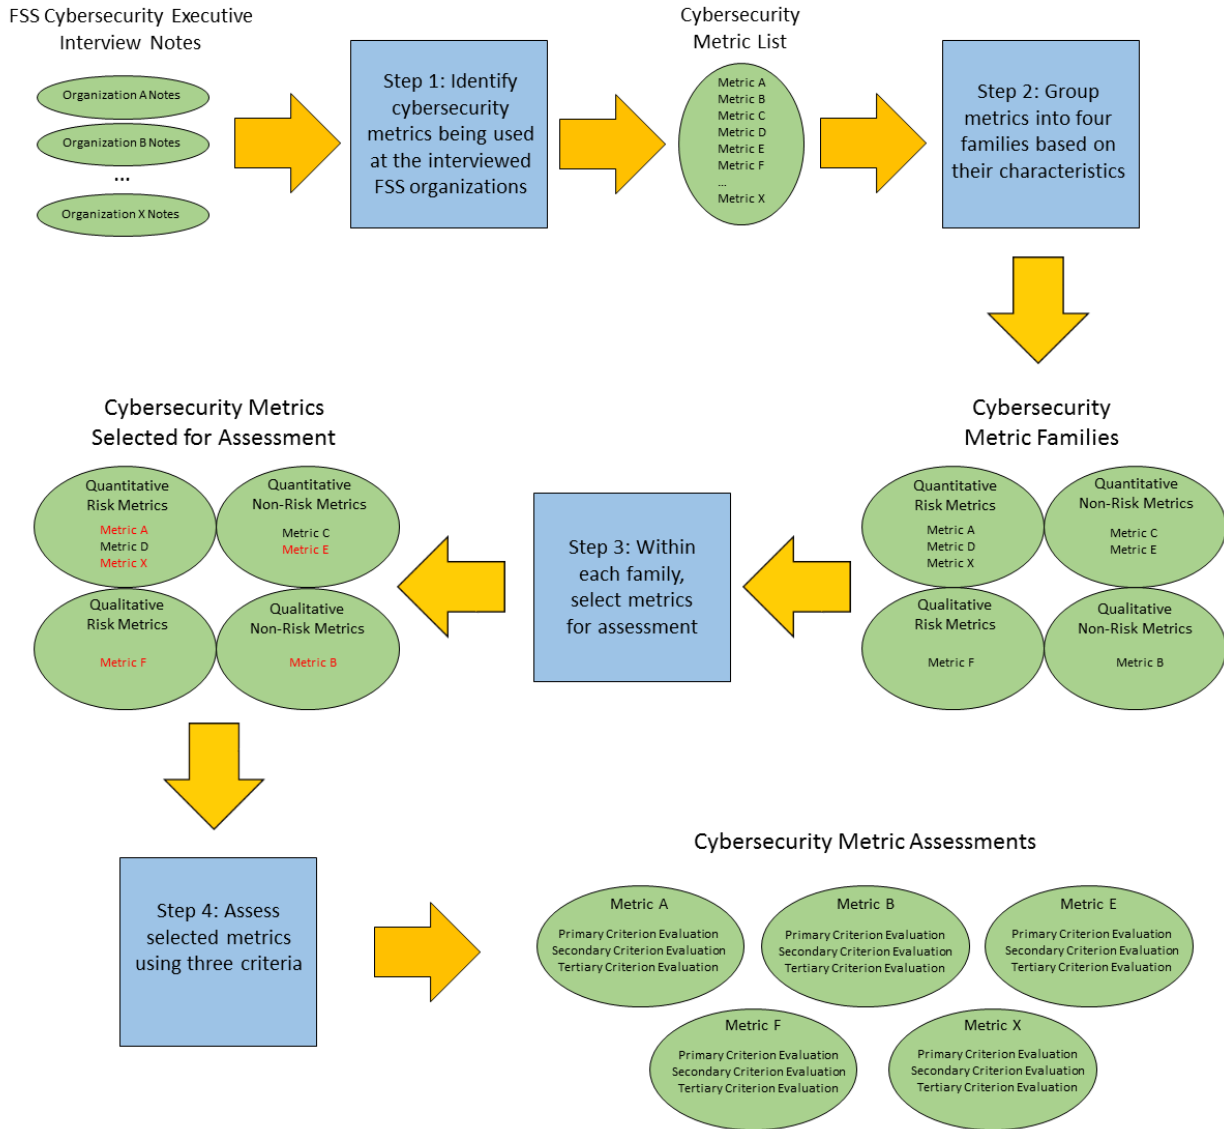


Figure 6. Cybersecurity Metric Assessment Methodology

6.3 Cybersecurity Metric Assessment Criteria

Within each cybersecurity metric family, a subset of metrics was selected for assessment against three criteria. These criteria were intended to determine the extent to which the metric benefits the originating organization and has the potential to benefit the NGCI program.

The primary criterion for evaluating a cybersecurity metric is the degree to which the metric is useful to the originating organization. As was indicated in section 2.2, the purpose of a risk

metric is to be useful. In the case of risk metrics, the importance of being useful even exceeds that of being accurate. This focus on usefulness is supported by the executive interviews which indicated that participating organizations seek usefulness and value from their efforts for regulatory compliance. During the interviews, it was also revealed that organizations often view metrics as having lifecycles where metrics are conceived to address a specific need and are retired once their usefulness decreases. As it is costly to compute metrics, the funding of a metric over time is contingent on its proven usefulness over time. Therefore, usefulness serves as a meaningful criterion for evaluating the benefit of a metric to its originating organization.

The secondary criterion for evaluating a cybersecurity metric pertains to whether the metric is transparent to its originating organization and whether the originating organization is applying the metric correctly. When metrics are used to inform decision making, it is vital that they are interpreted correctly and their underlying assumptions are properly understood. It is also important that decision makers do not view metrics as a black box. Rather, they should ensure that they have a thorough understanding of a metric's limitations and the proper situations in which a metric can be applied. Specifically, it is vital that a metric only be applied in situations where the data sources required by the metric are available at a sufficient quality. When a metric lacks the transparency necessary to ensure: (a) it is properly interpreted, (b) its assumptions are understood, (c) and its limitations are recognized, this can result in a situation where decision makers feel a false sense of security which, in turn, can lead them to make risky decisions. Furthermore, when a metric is applied in an improper situation or when data availability and quality are not present, it can result in decision makers drawing incorrect conclusions. Thus, a metric's transparency and the extent to which a metric is being correctly applied serves as a meaningful criterion for evaluating the benefit of a metric to its originating organization.

The tertiary criterion for assessing a cybersecurity metric is the extent to which the metric can be mapped or extended to the sub-sector, sector, and inter-sector levels. The FSS is one of the most tightly coupled yet competitive sectors. The interconnectedness between FSS organizations means that cybersecurity incidents occurring at just a small number of organizations can affect the sector as a whole. However, due to the high level of competition, no one private organization has the incentive to secure all other financial institutions and market utilities. Therefore, measuring and attempting to reduce cybersecurity risk in the FSS is an inherently governmental role. This criterion is intended to evaluate metrics based on their ability to fulfill the needs of the NGCI program with regards to the measurement and reduction of cybersecurity risk at the levels of sub-sector, sector, and across sectors.

6.4 Cybersecurity Metric Assessment Findings

While conducting interviews with financial institutions and market utilities, we formulated four broad observations. First, each financial institution and market utility we interviewed has a program of risk metrics and measures developed largely for that specific organization. Second, financial institutions and market utilities tended to utilize hybrid risk management frameworks specific to the organization rather than adopt one of the published frameworks. Third, risk modeling generally incorporated subjective assessments of threats and vulnerabilities. There were some efforts at quantification around consequence. And fourth, financial institutions and market utilities generally recognize their dependence on qualitative approaches and the need to

share data to inform the development of more quantitative metrics and mature their risk management capabilities. Below, we present a more detailed analysis according to our Cybersecurity Metric Assessment Methodology described above.

Our Cybersecurity Metric Assessment drew upon our interviews of several FSS executives. Each cybersecurity metric discovered during our review of these interviews was extracted and assigned to the appropriate, cybersecurity metric family. Within each family, we used two criteria for selecting metrics for assessment: (1) those deemed particularly useful or (2) those metrics deemed representative of the metrics for that particular family. These selected metrics were compared against the three assessment criteria outlined in section 6.3. As part of this comparison, each combination of metric and criterion was assigned a color indicating the extent to which the metric satisfies the criterion. Table 1 outlines each color used in the assessment along with a corresponding description.

Table 1. Cybersecurity Metric Assessment Color Key

Color	Description
	The criterion is currently satisfied for one or more organizations.
	The criterion is currently not satisfied and major revisions to the metric or implementation would be required to satisfy the criterion.
	The criterion is currently not satisfied but could be satisfied in the future using improved data sources.
	Further research is required to determine whether the criterion is satisfied.

Because the FSS executive interviews were loosely structured discussions, it was determined that a quantitative assessment of the interview data was infeasible and therefore a qualitative assessment was most appropriate. However, it should be noted that qualitative assessments such as this are subject to a number of limitations. Namely, they are inherently subjective and lack transparency, repeatability, and reproducibility. These limitations, which are discussed in greater detail within section 7.1, should be considered when utilizing the results of the Cybersecurity Metric Assessment.

6.4.1 Quantitative Risk Metrics

Table 2 depicts the Quantitative Risk Metrics Assessment.

Table 2. Quantitative Risk Metrics Assessment

#	Metric	Primary Criterion	Secondary Criterion	Tertiary Criterion
1	Risk methodology mapped to Basel methodology.	Green	Yellow	Yellow
2	Risk = f(threat, likelihood, impact).	Green	Red	Red
3	Risk assessments for all new software products and hardware expressed as a dollar amount.	Green	Yellow	Red
4	Large number of metrics derived from the NIST framework (265 monthly metrics, 140 in development, 130 retired). Metrics displayed in a dashboard and trail reality by a month due to lag time. A subset of the metrics deal specifically with risk. This subset includes: risk at login, risk associated with end of life applications, risk rating for all applications, vendor management risk, etc.	Green	Yellow	Red
5	Key Risk Indicator (KRI) featured in Gartner research.	Green	Yellow	Yellow
6	Risk = vulnerability x threat x asset value x probability of occurrence. Formula is currently used as guidance. Organization is not yet at the point where it is completely quantifiable and is not using the calculation yet.	Green	Red	Red
7	Some organizations evaluate cybersecurity as part of Comprehensive Capital Analysis and Review (CCAR) stress testing. CCAR is an annual exercise conducted by the Federal Reserve. A strong cyber program could allow an organization to hold less capital.	Yellow	Yellow	Yellow
8	Currently preparing for rollout of Factor Analysis of Information Risk (FAIR).	Green	Blue	Blue

6.4.2 Quantitative Non-Risk Metrics

Table 3 depicts the Quantitative Non-Risk Metrics Assessment.

Table 3. Quantitative Non-Risk Metrics Assessment

#	Metric	Primary Criterion	Secondary Criterion	Tertiary Criterion
1	Developing quantitative benchmarks from penetration testing and compliance audits.	Green	Yellow	Red
2	Budget percent allocated to protect against breach.	Green	Green	Red
3	Metrics specifically tailored to a problem (e.g., 90% reduction in misdirected messages).	Green	Yellow	Red
4	Rating controls on a scale from 1 to 10. Based on a capability maturity approach using categories such as: in planning, partial, full, and optimized.	Green	Red	Red
5	Measures investigations, security incidents, and end user reported issues.	Green	Green	Red
6	Determine cybersecurity spending by comparable firms and compare it to the spending carried out by the originating organization.	Green	Green	Red
7	Cyber-attack detection probability quantified using third party automated red teams. Organization is currently looking into this technology. It is not yet implemented.	Blue	Blue	Red
8	Large number of metrics derived from the NIST framework (265 monthly metrics, 140 in development, 130 retired). Metrics displayed in a dashboard and trail reality by a month due to lag time. A subset of the metrics is non-risk. This subset includes: impact likelihood similar to VCG matrix, vulnerabilities against service-level agreements by environment and business unit, intelligence alerts, etc.	Green	Yellow	Red
9	Key Performance Indicator (KPI) featured in Gartner research.	Green	Yellow	Red
10	Probability of occurrence for an attack by an external threat.	Green	Yellow	Red
11	The Cyber Defense Matrix (U Matrix) contains 5 rows and 5 columns. Rows include devices, applications, network, data, and people. Columns include identify, protect, detect, respond, and recover. The items included in matrix cells can vary based on the specific application of the matrix. In some applications, quantitative values such as expenditure, effective half-life, defense-in-depth score, defense-in-depth score combined with effective half-life, etc. can be represented.	Green	Yellow	Red
12	Regulators use dollars-per-headcount when evaluating budgeting for cybersecurity.	Yellow	Green	Red

6.4.3 Qualitative Risk Metrics

Table 4 depicts the Qualitative Risk Metrics Assessment.

Table 4. Qualitative Risk Metrics Assessment

#	Metric	Primary Criterion	Secondary Criterion	Tertiary Criterion
1	Security risk registry that classifies risk based on significance (minor, real, out of business).	Green	Yellow	Red
2	5x5 matrix that covers likelihood and impact ("Routing to" of various risks and includes a variety of qualitative assessments such as high/medium/low, unlikely to almost certain, and "Routing to" escalated).	Green	Yellow	Red

6.4.4 Qualitative Non-Risk Metrics

Table 5 depicts the Qualitative Non-Risk Metrics Assessment.

Table 5. Qualitative Non-Risk Metrics Assessment

#	Metric	Primary Criterion	Secondary Criterion	Tertiary Criterion
1	Control heat mapping used for board reporting. Heat map covers 11 to 14 streams (where they compare to industry, financial services, world class, etc.). Heat map shows where they can be if they invest certain amounts.	Green	Yellow	Red
2	Baseball cards for threat actors describing characteristics, organization experience with the threat actor, and industry experience with the threat actor.	Green	Green	Red
3	Organization has an automated system that looks at asset health for the top applications by dollar amount.	Green	Yellow	Red
4	The Cyber Defense Matrix (U Matrix) contains 5 rows and 5 columns. Rows include devices, applications, network, data, and people. Columns include identify, protect, detect, respond, and recover. The items included in matrix cells can vary based on the specific application of the matrix. In some applications, qualitative concepts such as market segments, security technologies, functional capability, courses of action, standards, etc. can be represented.	Green	Green	Red

6.5 Synthesis for the NGCI Apex Program

Not surprisingly, in Tables 2-5 most metrics met Criterion 1 for usefulness. In addition to describing the usefulness for a given metric, some interviewees also described their approach for assessing the usefulness of the metric itself. In some cases, they also described their approach for sunseting or retiring a given metric once its usefulness had diminished beneath a specified threshold.

Tables 2-5 also include a few instances of metrics which met Criterion 2 for transparency. In these instances, it was determined that it was unlikely the metric would be misunderstood within

the originating organization. Further, it was determined that the originating organization would likely have access to data necessary to compute the metric.

It is important to highlight the absence of any metric currently meeting Criterion 3 for scalability. This absence is largely driven by a lack of metrics that can currently be used to quantify cybersecurity risk beyond the organization level.

7 Quantitative Approach to Cybersecurity Risk Metrics

In this section, we recommend that the NGCI Apex program pursue the development of a quantitative cybersecurity risk metrics framework for the FSS. We begin with a brief review of the arguments for and against a quantitative approach to cybersecurity risk metrics. We then review two promising approaches to quantitative, cybersecurity risk metrics. Lastly, we conclude this section with an overview of initial steps for a confidence survey to support the NGCI program.

7.1 Argument in Favor of a Quantitative Approach to a Cybersecurity Risk Metrics Program

There are at least four aspects to the argument in favor of a quantitative approach to a Cybersecurity Risk Metrics program: transparency, repeatability, reproducibility, and scalability. These aspects of a quantitative risk metrics program extend beyond the limitations of qualitative assessments, to include the one above and similar, qualitative approaches.

7.1.1 Transparency

Consistent with Criterion 2 from the Assessment section above, *Transparency* reflects the antithesis of the “Black Box.” Transparency reveals the transformation from input data to output data without any ambiguity. This includes the full disclosure of all guiding assumptions as well as delimiting conditions.

7.1.2 Repeatability

Repeatability depicts the characterization of yielding a small variation in measurements when undertaken by a single researcher or instrument on the same item, and under the same conditions over a short period of time. Repeatability conveys consistency in the transformations between inputs and outputs.

7.1.3 Reproducibility

A hallmark of the scientific method, *Reproducibility* depicts the characterization of yielding a small variation in measurements when duplicating an entire experiment or study, typically by an independent researcher. Reproducibility serves as independent confirmation of the consistency in the transformations between inputs and outputs.

7.1.4 Scalability

Consistent with Criterion 3 from the Assessment section above, *Scalability* depicts the characterization of the efficient and transparent application of a metric from the level of an individual institution to aggregate levels of the Financial Services Sector (e.g., sub-sector and sector levels) and between sectors. A scalable approach allows for repeatable and reproducible aggregations of institutional metrics.

7.2 Argument Against a Quantitative Approach to a Cybersecurity Risk Metrics Program

There are at least two aspects to the argument against a quantitative approach to a Cybersecurity Risk Metrics program: a fundamental reliance on the quality and quantity of input data and potential confusion with precision. Below, we describe these aspects of the argument against a quantitative approach to a Cybersecurity Risk Metrics program. We then describe our recommendations to mitigate these aspects.

7.2.1 Reliance on Input Data

Any analytical approach to a metrics program, whether qualitative or quantitative, has a fundamental reliance on the quality and quantity of the input data which feeds the metrics program. The age-old adage of “Garbage In, Garbage Out” certainly applies here with a quantitative approach. A robust, risk metrics program relies on rich datasets – comprehensive data that has been collected and curated in a structured, and therefore repeatable, manner.

7.2.2 Confusion with Precision

Precision is widely considered to be the repeatability and reproducibility of a series of measurements within a sample (i.e., consistency within the measurements); whereas accuracy addresses the potential deviation of a series of measurements from the known, true value (i.e., consistency between the measurements and the actual value). Precision and accuracy may be conflated. While the quantity of data might not exist to provide a precise, point prediction of risk, the same data might suffice to give an accurate range of risk. Data only needs to be as precise as necessary to make an informed decision, which is the intrinsic purpose of the metric.

Advanced, quantitative methods should not convey a false degree of precision, a degree of precision beyond the precision of the raw data itself. Related to the points above regarding reliance on input data, quantitative methods cannot transform the precision of the input data. Therefore, quantitative methods should not be used to convey any impression of increased precision.

7.2.3 Mitigations

The following steps mitigate these aspects of the argument against a quantitative approach to a Cybersecurity Risk Metrics Program:

- Transparency in the transformation from input data to output data ensures that the quality and quantity of the input data are known to all stakeholders throughout the decision-making process.
- The Cybersecurity Risk Metrics Program should emphasize metrics with an accurate assessment of the range of risk, rather than metrics with a precise, point prediction.

7.3 Candidate Quantitative Frameworks for Implementation

The use of qualitative risk assessments, such as the standard risk matrix (Figure 7), to support cybersecurity control investment decisions can present numerous challenges. These challenges include the inherent lack of Transparency, Repeatability, Reproducibility, and Scalability present in these often subjective evaluations of risk. Furthermore, qualitative risk assessments can make it difficult to measure and visualize the marginal reduction in risk resulting from a proposed investment in one or more cybersecurity controls. One way to avoid these challenges is to base investment decisions off quantitative models evaluating cybersecurity risk.

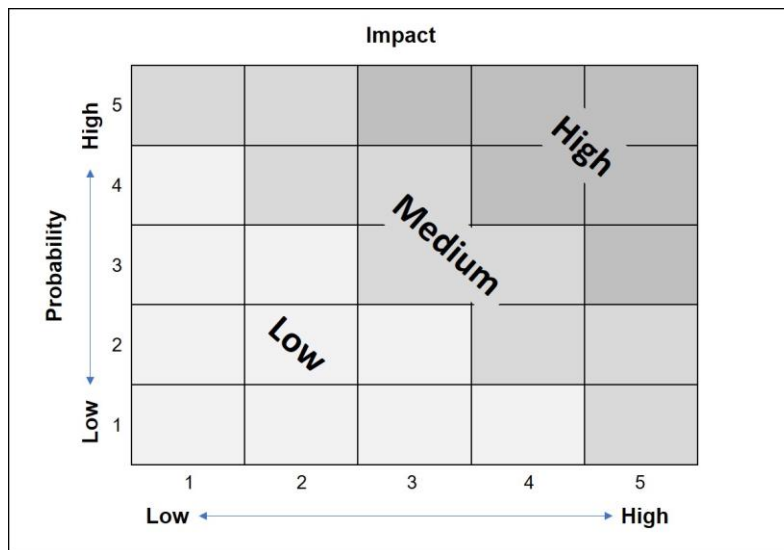


Figure 7. Standard Risk Matrix

Two prominent quantitative models assessing cybersecurity risk are FAIR [Josey 2014], which was introduced in section 3.2.2, and a methodology outlined by Douglas Hubbard and Richard Seiersen in the book *How To Measure Anything In Cybersecurity Risk* [Hubbard 2016]. Both FAIR and the Hubbard approach strive to create mathematically justifiable methods for evaluating cybersecurity risk within an organization. Therefore, both of these models have the potential to quantify cybersecurity risk within financial institutions and market utilities in a manner that could guide the NGCI program’s investment in controls. Below, we describe each of these two paths to achieving quantitative cybersecurity risk metrics and provide comparisons between them.

7.3.1 FAIR

As was discussed in section 3.2.2, FAIR utilizes the Open Risk Taxonomy to organize the quantitative components that allow for a computation of information security risk. Within the taxonomy, risk is defined as a function of loss event frequency (e.g., 0.24 loss events per year) and loss magnitude (e.g., \$120,000 per loss event). Loss event frequency and loss magnitude are each defined as a function of multiple components many of which are further divided into additional components. Therefore, the FAIR Open Risk Taxonomy allows an analyst to estimate

values for components of risk and then assemble those values into an overall evaluation for the risk within an organization.

Furthermore, FAIR allows for a quantitative analysis of alternatives to support decision makers attempting to choose between multiple cybersecurity controls in a resource constrained environment. Because the FAIR risk analysis method utilizes stochastic techniques such as Monte-Carlo simulation, it allows for an ultimate picture of organizational risk that incorporates the uncertainty present within the components of risk.

Although FAIR serves as a promising model for quantifying cybersecurity risk within financial institutions and market utilities, it also has a number of limitations. In order to compute risk by combining the components of risk, an analyst must first be able to obtain quantitative representations for the components of risk. This task would likely be challenging given the lack of historical data available within the financial sector to populate these components. It is possible that further analysis will reveal some sources of data that can be used to inform the risk components. However, it is likely that at least some components will need to be informed using expert opinion which can reduce the accuracy and increase the uncertainty surrounding the estimated values.

7.3.2 Hubbard and Seiersen Approach

In a manner similar to FAIR, the Hubbard and Seiersen (H&S) approach relies heavily on loss event frequency and loss magnitude to quantify cybersecurity risk. In order to implement the H&S approach, an analyst would identify a time horizon for the study (e.g., one year) as well as a list of risk events that could occur during the time horizon. For each risk event, historical data or subject matter experts are used to estimate the probability that the risk event occurs during the time horizon and a 90% confidence interval for the monetary loss if the risk event were to occur. Using the 90% confidence intervals, a lognormal distribution is created for each risk event representing the monetary loss due to an occurrence of the risk event. The list of risk events, with associated probabilities of occurrence and lognormal distributions are then used as inputs to a Monte-Carlo simulation analysis [Hubbard 2016].

Within this analysis, a single scenario representing the specified time horizon considers each risk event. For each risk event, random number draws are used to determine, based on the probability of occurrence, whether the event occurred in the scenario. If an event occurs, a random draw from the event's lognormal distribution is used to determine the monetary loss for the event. If the event does not occur, the monetary loss is \$0.

For the scenario, the monetary losses for each risk event within the scenario are summed to yield the total monetary loss. Because each scenario represents the specified time horizon, the monetary loss can be interpreted as the loss accrued during a simulated period equivalent to the time horizon. Therefore, if the time horizon for the study was 1 year and a scenario generated a total monetary loss of \$1 Million, the scenario represents a simulated year in which \$1 Million was lost due to the occurrence of risk events.

By simulating a large number of scenarios (e.g., 10,000 scenarios), an analyst can develop an understanding of the monetary loss distribution for the specified time horizon. This understanding is then used to create a Loss Exceedance Curve (LEC) chart indicating how, for

the specified time horizon, the “Chance of Loss or Greater” varies based on monetary loss amount [Hubbard 2016]. In order to analyze the reduction in risk due to proposed additional controls, additional LECs can be added to the chart to represent risk before the application of the proposed controls (inherent risk) and after the application of the proposed controls (residual risk).

Overall, the LEC chart is an effective tool for understanding and communicating quantitative risk. Rather than conveying simply a point estimate for the expected monetary loss, LEC charts provide a decision maker with a visual representation for the amount of uncertainty surrounding monetary loss.

Despite this strong advantage, the H&S approach also has limitations that mirror those of FAIR. Specifically, in the absence of historical data, the methodology would require expert opinion to populate its inputs. Although the H&S approach includes elicitation and calibration techniques designed to improve the quality of expert feedback, the use of expert opinion has the potential to cause inaccuracies in the input data that could alter the conclusions drawn through the analysis.

7.3.3 Comparison Between These Two Approaches

In general, both FAIR and the H&S approach appear to be consistent in that they both define cybersecurity risk as a function of loss event frequency and loss magnitude. Both approaches also attempt to quantify risk starting at the risk component level, representing the components quantitatively, and then using those quantitative representations to compute overall risk. Furthermore, both approaches allow an analyst to conduct an analysis of alternatives on proposed cybersecurity controls to determine the reduction in risk resulting from implementation of the controls. Interestingly, in their book, Hubbard and Seiersen [Hubbard 2016, p. 53] state that “In the authors’ opinion, FAIR, as another Monte Carlo-based solution with its own variation on how to decompose risk into further components, could be a step in the right direction for your firm.” At the very least, this endorsement highlights the philosophical consistencies between the two approaches while acknowledging some practical differences in the decomposition of risk.

7.4 Cybersecurity Situational Awareness Through a Confidence Survey

Due to the complexity of the cybersecurity landscape within the FSS, the development and implementation of a quantitative cybersecurity risk modeling framework for the sector will require a significant investment. Prior to the rollout of such a framework, the NGCI Apex program will require continuous feedback from the FSS to improve its situational awareness of cybersecurity issues within the sector. In order to gather this feedback, the program needs to administer a Confidence Survey to CART members. In the near term, this Confidence Survey will improve the program’s understanding of the overall cybersecurity risk present in the CART, the success of its past technology investments, and the CART’s desire for future technology investment. In the longer term, it will also gather data that can ultimately be used to inform a quantitative cybersecurity risk modeling framework for the sector. To this end, we have defined a set of eight characteristics deemed desirable in a Confidence Survey administered to CART members. The Confidence Survey:

1. Should be administered repeatedly over time at regular intervals.

2. Should produce a metric that quantifies cybersecurity risk in the CART.
3. Should produce metrics that can quantify loss event frequency in a future cybersecurity risk modeling framework for the FSS.
4. Should produce metrics that can quantify loss magnitude in a future cybersecurity risk modeling framework for the FSS.
5. Should be administered to all CART members.
6. Should not excessively burden its respondents.
7. Should produce metrics indicating how past NGCI Apex program investments in different technology areas have impacted cybersecurity risk in the CART.
8. Should produce metrics indicating the level of desire within the CART for future NGCI Apex program investment in different technology areas.

In section 8, we introduce four surveys developed by industry and academia that solicit feedback on cybersecurity issues. Then, in section 9.2, we derive six criteria that evaluate the extent to which an existing survey possesses the characteristics outlined above. These criteria serve as the basis for an assessment evaluating the ability of each survey introduced in section 8 to address the needs of the NGCI Apex program. Based on the results of the assessment, we provide a recommendation for how the NGCI Apex program should implement the Confidence Survey. Once the Confidence Survey is in place, NGCI stakeholders will be able to provide feedback throughout the lifecycle of the NGCI Apex program.

8 Introduction to Confidence Survey Exemplars

Within industry and academia, numerous surveys with associated analytic techniques are used to examine issues related to cybersecurity. These surveys tend to differ in their ability to fulfill the needs of the NGCI Apex program. In order to identify existing surveys that can fill the role of the Confidence Survey or contribute concepts useful when designing the Confidence Survey, we conducted a review of surveys developed by industry and academia. It should be noted that this was not intended to be an exhaustive review of all surveys examining cybersecurity related issues. Rather, it was intended to provide insights into concepts and techniques utilized in a subset of surveys. Table 6 provides a list of cybersecurity surveys included in the review. Subsequently, we briefly introduce each of these surveys.

Table 6. Cybersecurity Surveys Included in Review

#	Survey Name	Originating Organization(s)	Results Publication Release Interval	First Results Publication	Most Recent Results Publication
1	Index of Cyber Security	Dan Geer and Mukul Pareek	Monthly	April 2011	January 2017
2	Global Cybersecurity Status Report	ISACA and Cybersecurity Nexus (CSX)	Not Applicable	2015	2015
3	Cost of Data Breach Study: United States	Ponemon Institute and IBM	Annually	2005	2016
4	US State of Cybercrime Survey	PricewaterhouseCoopers, CSO, Carnegie Mellon University - Software Engineering Institute - CERT Division, United States Secret Service	Annually	2013	2015

8.1 Index of Cyber Security

The Index of Cyber Security (ICS) is a quantitative risk metric computed using data from a monthly survey administered to working cybersecurity experts across multiple industries. The Index, which is co-published by Dan Geer and Mukul Pareek, has been updated and released each month since its initial publication in April 2011. Formally, ICS is described as “A measure of perceived risk” where “A higher index value indicates a perception of increasing risk [and] a lower index value indicates the opposite” [Geer].

According to Geer and Pareek, one of the primary motivations for representing cybersecurity risk as an index is to enable the creation of financial instruments based on the index. In the future, organizations exposed to cybersecurity risk will likely hedge that risk by purchasing insurance. It is believed that the companies selling this insurance “will buy the derivatives [of ICS] because it is the insurance company that will be exposed to the “average” risk represented by the index” [Geer]. The derivatives will be sold by other entities within the financial market. Geer and Pareek note that the development of financial instruments based on ICS would be a complex

process. Therefore, the creation of these instruments and associated markets is considered outside the scope of their current work.

One of the greatest challenges faced when developing a quantitative cybersecurity risk metric is the inherent lack of data from which a metric can be computed. When developing ICS, Geer and Pareek encountered this challenge and ultimately concluded that a survey needed to be designed and repeatedly administered over time in order to establish a credible and reliable source of data for the Index. They note that potential sources of data such as prices on illegal markets, third party sources, or primary collection through interaction with individuals participating in illegal activity all have drawbacks that prevent them from being viable. Rather, they argue that utilizing a monthly survey as the data source for ICS provides: (1) “Coverage of a wider range of risks”, (2) “Better acceptance among industry”, (3) “Ease of maintenance”, (4) “Sub-indices on specific risks”, (5) “Practicality of implementation”, and (6) “Credibility” [Geer].

The survey contains questions that are generally consistent month to month. Geer and Pareek indicate that, although they desire for the question set to be consistent over time, infrequent modifications may occur. The initial version of the survey consisted of six sections covering the topics: (1) “Attack actors”, (2) “Weapons”, (3) “Effect desired by attackers”, (4) “Attack targets”, (5) “Defenses”, and (6) “Overall perceptions” [Geer]. Each section contains multiple questions related to the corresponding topic. Questions within the survey are “responded to on a five-point scale, with the respondents rating the risks as having fallen fast, fallen, stayed static, risen or risen fast when compared to the previous month” [Geer].

The value of ICS for March 2011, the month prior to the first publication of the index, was set equal to 1000. In order to standardize the process for updating the ICS value during each subsequent month based on survey results, Geer and Pareek developed a repeatable procedure.

An assessment of the value that the Index of Cyber Security survey could provide to the NGCI program is included in section 9.3.

8.2 Cost of Data Breach Study: United States

In section 7.3.3, it is noted that, in general, FAIR and the H&S approach define cybersecurity risk as a function of loss event frequency and loss magnitude. One of the challenges when implementing these and similar techniques is the lack of historical data to inform the loss event frequency or loss magnitude risk components. The United States Cost of Data Breach Study aims to produce quantitative estimates for loss magnitude. It is possible that, in some risk computation techniques, these estimates could be used to inform input parameters. At the very least, the study provides a wealth of information to improve situational awareness regarding the loss magnitude component of risk.

The United States Cost of Data Breach Study is an annual publication produced by the Ponemon Institute. Currently, the study is also sponsored by IBM. Since its initial publication in 2005, the study has been estimating and tracking a number of metrics related to the costs associated with data breach incidents at companies within the United States. In order to inform these metrics, the study applies benchmarks to companies that experienced a data breach. Two key metrics included in the study are “The average per capita cost of data breach” and “The average total

organizational cost of data breach” [Ponemon Institute 2016]. Both of these metrics include direct costs and indirect costs.

In the 2016 release, the authors note that “By design, [they] do not include cases involving more than 100,000 compromised records because they are not indicative of data breaches incurred by most organizations” [Ponemon Institute 2016]. When companies are identified for inclusion in a release of the study, a benchmark is administered to each company using “a confidential and proprietary benchmark method” [Ponemon Institute 2016]. The 2016 release states that benchmarks did not rely on “Actual accounting information” [Ponemon Institute 2016]. Rather, they “Relied upon numerical estimation based on the knowledge and experience of each participant” [Ponemon Institute 2016].

An assessment of the value that the Cost of Data Breach Study: United States could provide to the NGCI program is included in section 9.3.

8.3 Global Cybersecurity Status Report

The Global Cybersecurity Status Report was released in 2015 by ISACA and Cybersecurity Nexus (CSX). The report was based on “A global survey of 3,439 business and IT professionals in 129 countries [that captured] real-time insights on cybersecurity attacks, skills shortages and proposals from US President Obama” [ISACA 2015]. Findings from the report were released in the form of data sheets where there was one sheet covering global results and multiple additional sheets covering different geographic areas of interest.

The United States data sheet of the 2015 Global Cybersecurity Status Report includes feedback from 1,211 respondents living in the United States [ISACA 2015]. Of these respondents, 21% identified as working in the Financial/Banking Industry [ISACA 2015]. The data sheet covers the results from 10 questions related to cybersecurity issues. For each question, it provides a breakdown of the possible responses by percentage.

An assessment of the value that the Global Cybersecurity Status Report could provide to the NGCI program is included in section 9.3.

8.4 U.S. State of Cybercrime Survey

The US State of Cybercrime Survey was an annual survey “Co-sponsored by [PricewaterhouseCoopers], CSO, the CERT Division of the Software Engineering Institute at Carnegie Mellon University, and the United States Secret Service” [PricewaterhouseCoopers 2015]. Our review of the US State of Cybercrime Survey was based on the PricewaterhouseCoopers (PwC) results publication titled US cybersecurity: Progress stalled – Key findings from the 2015 US State of Cybercrime Survey [PricewaterhouseCoopers 2015]. The 2015 US State of Cybercrime Survey results were based on “Responses from more than 500 executives of US businesses, law enforcement services, and government agencies” [PricewaterhouseCoopers 2015].

An assessment of the value that the US State of Cybercrime Survey could provide to the NGCI program is included in section 9.3.

9 Assessment of Confidence Survey Exemplars

The NGCI program needs a Confidence Survey that can be continuously administered to the CART in order to improve the program’s situational awareness on cybersecurity issues. In section 8, we introduced four surveys examining issues related to cybersecurity. These surveys, which were developed by industry and academia, tend to differ in a number of areas that impact their ability to fulfill the needs of the NGCI program. Namely, they differ in the types of aggregate cybersecurity metrics they produce, the frequency at which metrics are updated based on new survey feedback, and the nature of the survey instrument itself regarding required characteristics of the respondents and the level of effort required by respondents to complete the survey. In order to understand how the four surveys perform in each of these areas and, by extension, the ability of each survey to fulfill the needs of the NGCI program, an assessment has been conducted.

9.1 Confidence Survey Exemplar Assessment Methodology

An assessment methodology was developed to evaluate the ability of each survey to fulfill the needs of the NGCI program.

When reviewing each of the surveys, we determined that a given survey usually has multiple result publications. These result publications can take the form of data sheets, summary charts, or more formal written reports. For a given survey, multiple result publications are usually spread out over time. For example, a survey administered annually may produce a single result publication each year in the form of a written report. Additionally, it is possible for a survey to release multiple result publications at roughly the same time. For example, a survey may release two result publications each month in the form of a detailed report made available to survey respondents and a set of high level summary charts made available to the general public. In order to scope our analysis, each survey’s assessment was primarily based on the evaluation of a single result publication selected for that survey. For a given survey, this evaluated result publication was selected based on the recency of its release date and its ability to provide a comprehensive overview of the survey’s results. Table 7 indicates the evaluated result publication for each survey [Geer 2016, ISACA 2015, Ponemon Institute 2016, PricewaterhouseCoopers 2015].

Table 7. Evaluated Result Publications

#	Survey Name	Evaluated Result Publication
1	Index of Cyber Security	The Index of Cyber Security - Detailed Report, October 2016
2	Global Cybersecurity Status Report	2015 Global Cybersecurity Status Report - US Data
3	Cost of Data Breach Study: United States	2016 Cost of Data Breach Study: United States
4	US State of Cybercrime Survey	US cybersecurity: Progress stalled - Key findings from the 2015 US State of Cybercrime Survey

In order to ensure that the four surveys identified in section 8 were assessed using a consistent process, six criteria were defined. Detailed descriptions of these criteria are provided in section 9.2. The six criteria, some of which include sub-criteria, were designed to assess the extent to which each survey possesses characteristics that would be considered desirable in a Confidence Survey administered to CART members. At a high level, the criteria cover topics including how frequently survey results are released, whether survey results are still being released, whether the survey produces certain kinds of cybersecurity metrics and the extent to which these metrics are tracked over time, and whether it would be feasible to administer the survey to CART members. For a given survey, the assessments for criteria pertaining to the kinds of metrics produced by the survey and whether those metrics are tracked over time are based exclusively on an analysis of the survey’s evaluated result publication. Assessments for other criteria are based on conclusions drawn from more general research into the survey that often included sources beyond the evaluated results publication.

The results of the survey assessments were used to populate the Cybersecurity Survey Assessment Matrix. Each column in the matrix corresponds to one of the assessed surveys and each row corresponds to one of the assessment criteria or sub-criteria. Each cell in the matrix provides the assessment result for the cell’s combination of survey and criterion or sub-criterion. This matrix and a discussion of the assessment results are provided in section 9.3.

9.2 Confidence Survey Exemplar Assessment Criteria

In section 7.4, we presented a list of eight characteristics deemed desirable in a Confidence Survey administered to the CART. If an existing survey within industry or academia were implemented to gather CART member feedback, it is possible that the survey could fulfill each of the first six characteristics in the list. However, as the final two characteristics pertain specifically to NGCI program operations, it is not possible for an existing survey to fulfill these characteristics if administered to the CART. Therefore, in order to efficiently evaluate each section 8 survey’s ability to meet the needs of the program, six assessment criteria were derived

from the first six characteristics in the list. Below, we describe the six criteria and associated sub-criteria used in the assessment of cybersecurity surveys.

9.2.1 Criterion One

The first criterion pertains to whether the survey is administered repeatedly over time enabling the publication of survey results at regular intervals. If a survey's assessment for the first criterion is "Yes", the survey is evaluated against two additional sub-criteria. These sub-criteria inquire about the length of the interval and whether it is reasonable to assume that the survey will be administered in the future with subsequent publication of results. Once implemented, the Confidence Survey should be able to provide insight into how metrics representing cybersecurity risk, components of cybersecurity risk, program investment performance, and demand for future program investment have changed over time. In order to provide this insight, the Confidence Survey will need to be administered repeatedly, at regular intervals, allowing for trend analysis of cybersecurity metrics. Shorter intervals would provide more granularity in the trend analysis but could lower the survey response rate if the burden to complete the survey is high. Likewise, longer intervals would reduce the trend analysis granularity but could potentially increase the survey response rate. If a survey will be administered in the future with subsequent publication of results, this increases the likelihood that collaboration between the NGCI program and the survey's originating organization may be possible thus allowing the survey to be administered to CART members.

9.2.2 Criterion Two

The second criterion asks whether the survey's evaluated result publication includes one or more metrics quantifying cybersecurity risk. If a survey's assessment for the second criterion is "Yes", the survey is evaluated against one additional sub-criterion. This sub-criterion asks whether the publication indicates how at least one of these metrics has changed over time. As was noted above for the first criterion, the NGCI program needs to be able to track the cybersecurity risk present in the CART over time. The ability to show temporal trends in cybersecurity risk will help the program communicate the importance of continued investment towards mitigating the risk. Additionally, unexpected changes in cybersecurity risk can be used to prompt discussions with the CART. These discussions could provide the program with insight into factors driving the risk fluctuation, which could help the program better prioritize its technology investments. Therefore, it is important that a Confidence Survey administered to CART members can repeatedly and consistently gather the feedback needed to show cybersecurity risk trends over time.

9.2.3 Criterion Three

The third criterion asks whether the survey's evaluated result publication includes one or more metrics quantifying the loss event frequency component of cybersecurity risk. If a survey's assessment for the third criterion is "Yes", the survey is evaluated against one additional sub-criterion. This sub-criterion asks whether the publication indicates how at least one of these metrics has changed over time. Section 7.3.3 indicates that both FAIR and the H&S approach, two promising alternatives for an FSS cybersecurity risk modeling framework, generally define cybersecurity risk as a function of loss event frequency and loss magnitude. Therefore, if a

Confidence Survey administered to CART members is able to solicit feedback that can be used to quantify loss event frequency over time, this could be of benefit to the NGCI program's attempt to establish a quantitative cybersecurity risk modeling framework.

9.2.4 Criterion Four

The fourth criterion asks whether the survey's evaluated result publication includes one or more metrics quantifying the loss magnitude component of cybersecurity risk. If a survey's assessment for the fourth criterion is "Yes", the survey is evaluated against one additional sub-criterion. This sub-criterion asks whether the publication indicates how at least one of these metrics has changed over time. As was noted above for the third criterion, both FAIR and the H&S approach generally define cybersecurity risk as a function of loss event frequency and loss magnitude. Therefore, similar to the third criterion, if a Confidence Survey administered to CART members is able to solicit feedback that can be used to quantify loss magnitude over time, this could be beneficial to the NGCI program's attempt to establish a quantitative cybersecurity risk modeling framework.

9.2.5 Criterion Five

The fifth criterion pertains to whether, if the survey were administered to CART members, the survey would allow for substantive participation by all CART members. This criterion focuses on the requirements the survey places on its respondents and whether those requirements would exclude a subset of CART members from participating in the survey. For example, a survey could require that all of its respondents be representatives from organizations that have experienced a significant cyber attack in the past year. If this were the case, CART members who did not experience such an attack would be precluded from participating in the survey. If a subset of CART members is prevented from participating in a given survey, the NGCI program could still potentially benefit from administering the survey to those CART members eligible to participate. However, the program would need to ensure that there is a large enough group of respondents with sufficient organizational diversity to allow for the computation of metrics that are representative of the CART or FSS. Furthermore, it should be noted that limiting survey respondents to a specific subset of CART members could bias feedback due to similarities among respondents. If all respondents are from organizations that have recently experienced a large cyber attack, the feedback gathered using the survey may not be representative of all organizations in the CART or sector, especially if much of the feedback is subjective in nature.

9.2.6 Criterion Six

The sixth criterion asks whether the expected burden to complete the survey would be small enough to allow the survey to be administered to CART members. As was noted above for the fifth criterion, it is vital that a survey has a large enough group of respondents with sufficient organizational diversity to allow for the computation of metrics that are representative of the CART or FSS. In order to increase the likelihood that a survey administered to CART members will have enough respondents to provide valuable feedback, it is important to maximize the response rate among CART members who receive the survey. One critical component of maximizing the response rate is ensuring that the amount of time required to complete the survey is not so large that it discourages participation.

9.3 Confidence Survey Exemplar Assessment Results

In order to determine which surveys possess key characteristics deemed desirable in a Confidence Survey administered to CART members, an assessment was conducted. The results of this assessment were used to produce the Cybersecurity Survey Assessment Matrix, which is provided in Table 8.

Table 8. Cybersecurity Survey Assessment Matrix

		Cybersecurity Survey			
		Index of Cyber Security	Global Cybersecurity Status Report	Cost of Data Breach Study: United States	US State of Cybercrime Survey
Cybersecurity Survey Assessment Criterion	1) Is the survey administered repeatedly over time enabling the publication of survey results at regular intervals?	YES	NO	YES	YES
	a) If yes, what is the interval?	MONTHLY	NOT APPLICABLE	ANNUALLY	ANNUALLY
	b) If yes, is it reasonable to assume that the survey will be administered in the future with subsequent publication of results?	YES	NOT APPLICABLE	YES	NO
	2) Does the survey's evaluated result publication include one or more metrics quantifying cybersecurity risk?	YES	NO	NO	NO
	a) If yes, does the publication indicate how at least one of the metrics has changed over time?	YES	NOT APPLICABLE	NOT APPLICABLE	NOT APPLICABLE
	3) Does the survey's evaluated result publication include one or more metrics quantifying the loss event frequency component of cybersecurity risk?	NO	YES	NO	YES
	a) If yes, does the publication indicate how at least one of the metrics has changed over time?	NOT APPLICABLE	NO	NOT APPLICABLE	NO
	4) Does the survey's evaluated result publication include one or more metrics quantifying the loss magnitude component of cybersecurity risk?	NO	NO	YES	NO
	a) If yes, does the publication indicate how at least one of the metrics has changed over time?	NOT APPLICABLE	NOT APPLICABLE	YES	NOT APPLICABLE
	5) If the survey were administered to CART members, would the survey allow for substantive participation by all CART members?	YES	YES	NO	YES
	6) Would the expected burden to complete the survey be small enough to allow the survey to be administered to CART members?	YES	YES	NO	UNKNOWN

9.3.1 Assessment for Index of Cyber Security

The first survey reviewed during the assessment was the Index of Cyber Security (ICS). Overall, this survey has several characteristics that would allow it to provide value to the NGCI program. Of the four surveys reviewed during the assessment, ICS was the only survey that included a metric quantifying cybersecurity risk in its evaluated result publication. This metric, the Index of Cyber Security value, is updated monthly and will continue to be updated in the future. Because ICS was the only survey identified as producing a cybersecurity risk metric, it is the only assessed survey that offers the possibility of quantifying cybersecurity risk within the CART if administered to CART members. In addition to providing the Index of Cyber Security value, the evaluated result publication also included information for sub-indices tracking components of cybersecurity risk. These sub-indices would help the NGCI program understand and explain fluctuations in the Index of Cyber Security value over time. Because all CART members are eligible to participate in ICS and the burden to complete the survey is sufficiently small, it would likely be feasible to administer ICS to CART members. One noteworthy drawback of ICS is that it did not include metrics that quantify loss event frequency or loss magnitude in its evaluated result publication. This is because the ICS Value is not defined as a function of the loss event frequency and loss magnitude components of risk. Therefore, although ICS can quantify cybersecurity risk, it would likely be of limited value when attempting to develop a quantitative cybersecurity risk modeling framework in alignment with FAIR or the H&S approach.

9.3.2 Assessment for Global Cybersecurity Status Report

The second survey reviewed during the assessment was the Global Cybersecurity Status Report. Overall, this survey would likely be of limited value to the NGCI program. The Global Cybersecurity Status Report's evaluated result publication does not include metrics quantifying cybersecurity risk or loss magnitude. Although it does include metrics that quantify loss event frequency, these metrics simply provide the percentages of respondents who answered yes, no, and unsure to a question about whether the respondents expect their organization to be hit with a cyber attack in 2015 [ISACA 2015]. These metrics can potentially provide value in improving high level situational awareness regarding loss event frequency. However, they are likely insufficient to inform cybersecurity risk models requiring a rate of occurrence for loss events. All CART members would be eligible to participate in the Global Cybersecurity Status Report and the burden to complete the survey would not prevent the NGCI program from being able to administer the survey to the CART. However, as the survey was only administered in 2015, it is not administered repeatedly over time enabling the publication of survey results at regular intervals. This makes it infeasible to administer the Global Cybersecurity Status Report survey to CART members as a means of gathering continuous feedback.

9.3.3 Assessment for Cost of Data Breach Study: United States

The third survey reviewed during the assessment was the Cost of Data Breach Study: United States. Overall, this survey has the potential to be of value to the NGCI program's attempt to establish a quantitative cybersecurity risk modeling framework. The survey's evaluated results publication does not include any metrics quantifying cybersecurity risk. Therefore, if the survey were administered to CART members, it would be unable to quantify the cybersecurity risk

present in the CART. Although the survey does not include any metrics quantifying loss event frequency in its evaluated result publication, the publication does provide values for quantitative loss magnitude metrics over time. Two valuable loss magnitude metrics produced by the study are “The average per capita cost of data breach” and “The average total organizational cost of data breach” [Ponemon Institute 2016]. It is possible that these metrics could be used to inform input parameters for quantitative cybersecurity risk models that divide risk into loss magnitude and loss event frequency components. Therefore, this survey has the potential to further the NGCI program’s attempt to establish a quantitative cybersecurity risk modeling framework. The United States Cost of Data Breach Study is unique among assessed surveys in that it only solicits feedback from companies that have experienced a data breach meeting certain criteria. Therefore, if the NGCI program wanted to administer the survey to the CART, it is likely that only a subset of CART members would be qualified to participate. Furthermore, the survey requires interviews with “IT, compliance and information security practitioners who are knowledgeable about their organization’s data breach and the costs associated with resolving the breach” for each participating company [Ponemon Institute 2016]. As a result, the high burden imposed on companies responding to the survey would likely make it infeasible to administer the survey to CART members. Despite this limitation, the fact that survey results are published annually and the survey will likely continue to publish results in the future means that the metrics quantifying loss magnitude will continue to be updated at regular intervals. Although these metrics would not be based on CART member feedback, they could still be of use when modeling cybersecurity risk within the FSS.

9.3.4 Assessment for US State of Cybercrime Survey

The fourth survey reviewed during the assessment was the US State of Cybercrime Survey. Overall, it is likely that this survey would be of limited value to the NGCI program. The survey’s evaluated result publication does not provide any metrics quantifying cybersecurity risk. As has been noted previously, this means that if the survey were administered to CART members, it would be unable to quantify the cybersecurity risk present in the CART. Although the survey’s evaluated result publication does not include any metrics quantifying loss magnitude, it does include metrics quantifying loss event frequency. These metrics tend to focus on the percentages of respondents who detected security incidents generally or of various types in the past year [PricewaterhouseCoopers 2015]. Additionally, there were metrics that provided the number of detected incidents at large, medium, and small companies in the past year [PricewaterhouseCoopers 2015]. Although all of these metrics contribute to overall situational awareness regarding loss event frequency, it would be either challenging or impossible to convert these metrics into rates of occurrence for loss events that could inform input parameters in a future quantitative cybersecurity risk modeling framework for the FSS. All CART members would be eligible to participate in the US State of Cybercrime Survey. However, a lack of information regarding the survey mechanism has limited our ability to obtain a clear understanding of the expected burden required to complete the survey. Although results from the survey were historically published on an annual basis, it does not appear that results were published in 2016. This could be an indication that the survey will not be administered in the future, which would make administering the survey to CART members infeasible.

10 Implementation Plan for an NGCI Confidence Survey

In section 7.4, we defined eight characteristics deemed desirable for an NGCI Confidence Survey. The ability of the Confidence Survey to benefit the NGCI Apex program depends on the extent to which the survey possesses these characteristics. The Confidence Survey should be administered to all CART members repeatedly over time at regular intervals. The burden to complete the survey should be low enough to ensure a high response rate among CART members receiving the survey. In addition to these general characteristics, the Confidence Survey should also compute several metrics that will improve situational awareness in areas critical to the NGCI program. The survey should quantify cybersecurity risk in the CART. It should also produce quantitative metrics indicating how past program investments have impacted cybersecurity risk in the CART and the desire within the CART for future program investments. The Confidence Survey should also inform the development of a cybersecurity risk modeling framework for the FSS, presumably in alignment with FAIR or the H&S approach. Once a framework is developed, the survey should assist in quantifying the loss event frequency and loss magnitude components of cybersecurity risk.

In section 8, we identified four surveys, developed by industry and academia, that examine cybersecurity issues. These surveys consisted of: (1) The Index of Cyber Security, (2) The Global Cybersecurity Status Report, (3) The United States Cost of Data Breach Study, and (4) The US State of Cybercrime Survey. An assessment of these surveys, presented in section 9, revealed that none of the surveys possessed all eight characteristics deemed desirable in a Confidence Survey administered to the CART. Of the assessed surveys, the Index of Cyber Security (ICS) and the United States Cost of Data Breach Study are the only surveys with the potential to provide considerable value to the NGCI Apex program. The assessment concluded that it would likely be feasible to administer ICS to CART members. However, it would likely be infeasible to administer the United States Cost of Data Breach Study to the CART due to the high burden the survey places on its respondents. Therefore, we recommend that ICS be used as the basis for a Confidence Survey administered to CART members. To implement this recommendation, the NGCI program should direct that HSSEDI extend its collaborative relationship with the ICS co-publishers, Dan Geer and Mukul Pareek.

As was mentioned previously, ICS has the potential to provide considerable value to the NGCI Apex program but does not possess all the required characteristics of an NGCI Confidence Survey. In section 10.1, we describe the ways ICS could provide value to the program. Then, in sections 10.2 to 10.4, we discuss three areas in which ICS should be expanded to meet the requirements of the NGCI program.

10.1 Benefits of Implementing ICS

There are two primary factors that allow ICS to provide a strong foundation for the Confidence Survey. The first is that ICS possesses all the Confidence Survey's desired characteristics pertaining to how the survey should be administered. The second factor is that ICS has the capability to quantify the perception of cybersecurity risk among CART members.

10.1.1 Administering ICS to CART Members

Currently, ICS is administered monthly using a consistent question set to inform its metrics. Survey results are compiled into a detailed report for each month. It is likely that we would follow a similar procedure if the decision were made to administer ICS to the CART. Administering ICS to CART members monthly using a consistent question set to inform metrics would ensure that metrics produced by the survey are updated in a consistent manner with uniform temporal spacing between updates. This will allow us to conduct trend analysis showing how the values of metrics have changed over time. Delivering survey results to the NGCI Apex program in the form of a detailed monthly report would provide the program with continuous feedback from CART members. Moreover, the metrics included in the monthly reports will help the program maintain situational awareness on issues impacting its operation. Because ICS can be administered to all CART members and does not impose an unreasonable burden on its respondents, it increases our likelihood of receiving a large enough group of responses with sufficient organizational diversity to inform metrics that are representative of the CART or FSS.

10.1.2 Quantifying the Perception of Cybersecurity Risk Among CART Members

As was indicated in section 9.3.1, ICS currently produces a quantitative cybersecurity risk metric referred to as the ICS value. The co-publishers of ICS describe the ICS value as being “A measure of perceived risk” where “A higher index value indicates a perception of increasing risk [and] a lower index value indicates the opposite” [Geer]. The ICS value is updated each month using data collected through 25 survey questions. In addition to informing the overall ICS value, data collected through a given question is also used to produce a sub-index for that question. This results in a total of 25 sub-indices describing factors that contribute to the ICS value.

If ICS is administered to the CART, we will likely compute the ICS value and 25 sub-indices in a manner similar to that described above. Thus, the NGCI Apex program will receive a quantitative metric, updated monthly, representing the perception of cybersecurity risk among CART members. This will provide value by ensuring that the program has continuous situational awareness regarding the CART’s collective perception of cybersecurity risk. This situational awareness may help the program justify the need for continued efforts to reduce cybersecurity risk in the FSS. In addition to an updated ICS value, the program will also receive updated values for each of the 25 sub-indices in its detailed monthly report. These sub-indices will help the program better understand the factors impacting the ICS value. For example, if a monthly report indicates an unexpected jump in the perception of cybersecurity risk, the sub-indices will provide insight into the specific survey questions that contributed most to the increase. This could help inform CART meeting discussions regarding the topics covered in the questions, further improving the program’s situational awareness and perhaps identifying ways to mitigate factors increasing the perception of cybersecurity risk.

As is noted above, the ICS value has the potential to benefit the NGCI Apex program in the short term by quantifying the perception of cybersecurity risk among CART members. In the longer term, the program will want to implement a permanent mechanism that estimates the cybersecurity risk present in the entire FSS. The ICS value’s reliance on subjective feedback to estimate the perception of cybersecurity risk lowers the desirability of using it as the basis for this permanent mechanism. As we will discuss in section 10.4, it is our opinion that a

cybersecurity risk modeling framework, leveraging concepts from FAIR or the H&S approach, would be a better long-term alternative for quantifying cybersecurity risk in the sector. Therefore, the ICS value should be viewed as a temporary metric improving the ability of the program to understand the perception of cybersecurity risk among CART members prior to the eventual rollout of an FSS cybersecurity risk modeling framework.

10.2 Expanding ICS to Quantify Impact of Past Investments

When organizations regularly engage in investing activities over time, they almost always desire a feedback mechanism to help them quantify the performance of their past investments. The motivation to acquire such a mechanism is twofold. First, organizations often need to explain and defend their past investment decisions to stakeholders. Metrics, produced by a feedback mechanism, that quantify past investment performance can be useful in facilitating this kind of stakeholder communication. A second motivation for obtaining a feedback mechanism is that, over time, metrics produced by the mechanism can help an organization understand factors that impact the performance of its investments. This understanding can then be used to improve the quality of the organization's future investment decisions.

Currently, the NGCI Apex program lacks a feedback mechanism quantifying the impact past technology investments have had on cybersecurity risk within the CART. This reduces the program's ability to effectively communicate its contributions towards improving cybersecurity in the FSS. Moreover, it limits the ability of the program to learn from the successes and shortcomings of its past investments. It is imperative that a Confidence Survey administered to CART members help fill this gap.

As was mentioned in section 10.1.2, ICS can quantify the perception of cybersecurity risk among CART members. In theory, knowledge of historical investment decisions and an understanding of how the perception of cybersecurity risk among CART members has changed over time could allow for trend analysis indicating how past investments have impacted the perception of cybersecurity risk. However, in practice, this approach would likely be unsuccessful in providing the program with useful feedback. There are many factors, beyond past investment decisions, that could impact the perception of cybersecurity risk among CART members. Additionally, there may be many different technology investments occurring in tandem over time. These two factors, coupled with the limited number of data points that would be produced by a monthly Confidence Survey, make it unlikely that this approach would yield metrics that could accurately quantify the impact of past investments on the perception of cybersecurity risk among CART members. Therefore, an alternate approach would be needed to provide the program with insight into the performance of its investments. We believe that this approach should be based on supplemental questions added to the Confidence Survey.

Further research will be needed to determine how best to produce metrics that quantify the impact of past investments using supplemental Confidence Survey questions. One potential way forward would be to employ a methodology similar to that currently used in ICS. We could define a set of technology areas such that each technology investment would be viewed as an investment in one area. For each technology area, a single index would be created representing the confidence of CART members that past program investments in the area have reduced the current cybersecurity risk in their organizations. Each index would be updated monthly using

feedback from a single confidence survey question and a procedure similar to that used when updating the ICS value. The values of a given index over time would be compared to a curve indicating the monetary investment in the index's technology area by month. This comparison could help the program understand and communicate the performance of past investments in the technology area. Discussing the results of these comparisons at CART meetings may provide insight into why investments in certain technology areas were more effective than investments in other areas. This insight could help improve the quality of future NGCI Apex program investment decisions.

10.3 Expanding ICS to Quantify Demand for Future Investments

For the NGCI Apex program to properly invest its limited monetary resources, it must have a clear understanding of the CART's desire for investment in different technology areas. This understanding would help the program align its investments with the needs of the CART, increasing the likelihood that technologies funded by the investments are ultimately implemented in the FSS. Given that the Confidence Survey will be administered to all CART members monthly, it provides the NGCI Apex program with an excellent opportunity to gather continuous feedback on the CART's demand for investment in different technology areas. In its current state, ICS will not produce metrics providing this kind of feedback. Therefore, supplemental questions should be added to the Confidence Survey to allow for the computation of such metrics.

Further research will be required to design a process that uses supplemental Confidence Survey questions to inform metrics quantifying the CART's desire for investment in various technology areas. However, it may be possible to base this process on the methodology currently used in ICS. Specifically, we could follow an approach similar to that discussed in section 10.2. We would define a set of technology areas, each having its own index. An index would represent the desire of CART members for current NGCI Apex program investment in the given technology area to reduce the future cybersecurity risk in their organizations. Each index would be updated monthly using feedback from a single Confidence Survey question. The update would follow a procedure derived from that used when updating the ICS value. Over time, these indices would improve the NGCI Apex program's situational awareness regarding the CART's demand for investments in different technology areas. This improved situational awareness could help the program guide its investment decisions and align its investments with the needs of the CART.

10.4 Expanding ICS to Inform Cybersecurity Risk Modeling

As was mentioned in section 9.3.1, it is unlikely that ICS will be of value in furthering the development of a cybersecurity risk modeling framework for the FSS. In sections 10.4.1 and 10.4.2, we discuss two reasons it is important for the Confidence Survey to further the development of a cybersecurity risk modeling framework for the sector. The first pertains to a limitation of the ICS value, mandating the need for a more accurate cybersecurity risk metric. The second pertains to the ability of a cybersecurity risk modeling framework to guide the program's technology investments by estimating the reduction in risk resulting from different investment alternatives. Building on this discussion, in section 10.4.3 we outline how ICS can be

expanded to further the development of a quantitative cybersecurity risk modeling framework for the FSS.

10.4.1 Need for an Improved Cybersecurity Risk Metric

Up to this point, we have taken a relatively liberal interpretation of what it means to quantify cybersecurity risk. In doing so, we have implied that because the ICS value quantifies the perception of cybersecurity risk among CART members, it quantifies the cybersecurity risk present within the CART. Although the ICS value and the cybersecurity risk present within the CART are almost certainly correlated, they are not the same metric. This is because, at its core, the ICS value is based on sentiment. Each CART member answers 25 questions some of which pertain to the experience of their organization, others pertain to cybersecurity issues more generally. Although it is possible that, for a given CART member, some answers are based on objective metrics, it is more likely that answers are based on the member's subjective opinions. The responses of each CART member are aggregated into the ICS value representing the perception of cybersecurity risk among members of the CART.

Now, let us compare this approach to one that estimates cybersecurity risk rather than the perception of risk. Consider an instance where a future cybersecurity risk modeling framework is used to compute the aggregate cybersecurity risk present in the FSS. Such a framework would likely use loss event frequency and loss magnitude inputs to compute the cybersecurity risk present in individual organizations within the sector, either in the form of an expected rate of monetary loss or some other value. Cybersecurity risks produced for organizations could then be aggregated into an overall metric representing the cybersecurity risk present within the FSS.

It is evident that the metric produced using the cybersecurity risk modeling framework fundamentally differs from the ICS value in that it attempts to estimate cybersecurity risk rather than simply quantifying the perception of risk among members of the target population. As the primary goal of the NGCI Apex program is to reduce cybersecurity risk in the sector, a cybersecurity risk modeling framework for the FSS will likely produce metrics that estimate the value the program seeks to minimize more accurately than an ICS value modified to represent the full sector. Therefore, it is important ICS be expanded to help further the development of such a framework.

10.4.2 Ability to Evaluate Different Investment Alternatives

One advantage of using a model, rather than a sentiment-based survey, to estimate cybersecurity risk is that models often provide a deeper understanding of how risk components interact to drive overall cybersecurity risk at the organization and sector levels. This deeper understanding of the underlying system driving cybersecurity risk opens the door to analysis predicting how intentional modifications to this system could reduce risk. In section 7.3.3, we indicated that both FAIR and the H&S approach allow an analyst to conduct an analysis of alternatives on proposed cybersecurity controls to determine the reduction in risk resulting from implementation of the controls. This kind of predictive analysis would be either extremely challenging or impossible in instances where cybersecurity risk is quantified without an underlying model, as is the case with the ICS value. If the NGCI Apex program can eventually implement a cybersecurity risk modeling framework in the FSS, this could allow the program to predict the reduction in

cybersecurity risk at the organization or sector levels resulting from different technology investment decisions. This would provide value by helping guide the program's investments and providing the program with a means to justify its investment decisions.

10.4.3 Furthering Development of Cybersecurity Risk Modeling Framework

The process of implementing a cybersecurity risk modeling framework in the FSS will be challenging from a technical perspective. However, the greatest challenge will likely result from the large number of organizations, with differing priorities, that will need to collaborate for a framework to be operationalized. The NGCI Apex program is not affiliated with a regulatory agency. Therefore, its role is not to mandate that financial institutions adopt a certain approach to quantifying their cybersecurity risk. Rather, because of its unique relationship with the sector, the program is positioned to help build a consensus among financial institutions regarding a cybersecurity risk modeling framework. For such a framework to succeed, it will need to allow institutions the ability to compute their own organizational cybersecurity risk following a model tailored to their individual institution, subject to a set of constraints. These constraints should ensure that each institution's model allows for the computation of standard metrics that can be aggregated across all institutions in the FSS to estimate the overall cybersecurity risk in the sector. Our initial research suggests that such a framework may involve individual institutions computing their risk based on the loss magnitude and loss frequency of various cybersecurity events impacting the institutions. However, the eventual design for a framework would need to be developed and agreed upon by institutions in the FSS. Our belief is that the NGCI Confidence Survey could be useful in starting and facilitating this process.

Shortly after the Confidence Survey is implemented, we plan to begin incorporating supplemental questions into the survey to support the development of a cybersecurity risk modeling framework. These questions will initially be designed to further the program's understanding of how CART members currently estimate cybersecurity risk and the desire of CART members to move towards using quantitative models to estimate their risk. The results of these questions can be used to start CART meeting discussions about the importance of establishing a cybersecurity modeling framework in the sector and the benefits this framework will have on individual institutions and the sector as a whole. Assuming sufficient progress is made in generating support for a framework, the Confidence Survey can then be tailored to help guide CART member discussions regarding the technical details of the framework. When the design of the framework is complete, the Confidence Survey can gather data useful in producing input parameters for CART member organizational models. Presumably, these input parameters would relate to the loss magnitude and loss frequency of cybersecurity events occurring at institutions in the sector.

11 Conclusions and Recommendations

This report details a number of risk metrics and modeling frameworks from government, academia, standards bodies, and industry sources. The report also describes the findings and assessments from a series of interviews conducted with senior representatives from several financial sector critical infrastructure institutions. The *Implementation Plan* in section 10 synthesized the previous sections into an approach for the NGCI program to achieve a comprehensive Cybersecurity Risk Metrics program. Drawing from the *Implementation Plan*, HSSEDI makes the following three recommendations:

- The NGCI program should develop a scalable framework for cybersecurity risk metrics, drawing on concepts from two prominent modeling approaches.
- The NGCI program should implement a Confidence Survey to continuously gather feedback regarding cybersecurity issues that impact the CART. Over time, the feedback gathered from this survey could be used to guide and evaluate investment decisions for the NGCI program.
- In addition to being informative in its own right, the Confidence Survey should serve as an initial step on the trajectory to a scalable cybersecurity risk metric which meets the needs of the NGCI program. The NGCI program should evolve its Cybersecurity Risk Metrics program to extend beyond the Confidence Survey. The Cybersecurity Risk Metrics program should encompass a mature risk management program at the enterprise level and beyond, to include assessments of cybersecurity investments.

Appendix A Methodology for Survey of Metrics

The survey gathered metric information from three sources:

- Risk metrics and risk modeling frameworks promulgated by the government, security community, or financial sector
- Published literature on cyber risk metrics
- Interviews with CART members and other financial sector critical infrastructure institutions and SMEs.

Interviews with senior cybersecurity representatives responsible for threat modeling, risk assessment, and mitigation were generally conducted as one-hour phone conferences. They were held on a non-attribution basis to allow candid discussion without release of sensitive information about individual institutions, and indeed, participants were quite forthcoming and provided many valuable insights. The interviews were structured according to the following questions, to stimulate open-ended discussion.

1. What is the process you use to conduct threat modeling, understand and measure cyber risk, and prioritize investments to mitigate? Basically, how do you decide where to spend your next security dollar?
 - What are the components of your threat modeling and risk measurement approaches?
 - How do you identify cyber threats and map them to business functions?
 - What do you use to quantify elements of risk and impact in your executive/board reporting – value at risk quantification, professional judgement, qualitative criteria (e.g., FICO like scores, high/medium/low assessments)?
 - Do you have a Top 10 list – how do you define what to include, or decide if it is comprehensive enough?
 - What do you include from your vendor risk management assessment process?
 - Do you incorporate ‘end of life’ considerations or metrics?
 - How many staff work with you and how is the group structured?
2. Where/how do you get the data you need for the inputs to your modeling and measurement approaches? What data, if any, do you wish was available or would be helpful?
3. What are the outputs of your threat modeling and risk measurement approaches (e.g., dollar amount, heat map, trend lines, sentiment indicators)?
4. What is the feedback loop on the risk metrics or measurements you use – what is most convincing to your executive management (e.g., credibility, benchmarking)?
5. What can be extrapolated from your enterprise to the sub-sector, or sector level? Of the high priority threats/risks that keep you up at night, are they specific to your organization, systemic, or shared?

6. What, if anything, do you think would improve the sector's existing threat modeling or risk measurement approaches? If so, what and why?
7. At the sector level, how should government investments be selected and assessed or measured? Is there a scenario at the sector level that deserves more consideration, research, or investment than it is getting?
8. Are there any questions you think should be added to what we have asked, or that you would be interested in hearing responses to from the other organizations we will be interviewing?

The interview questions addressed the institutions' processes for both threat modeling and risk assessment. Findings germane to risk metrics are summarized in Section 4 of this report. Findings related specifically to threat modeling are provided in the companion Threat Model Survey and Assessment report.

List of Acronyms

Acronym	Definition
AFCEA	Armed Forces Communications and Electronics Association
BCBS	Basel Committee on Banking Supervision
CART	Cyber Apex Review Team
CAT	Cybersecurity Assessment Tool
CFPB	Consumer Financial Protection Bureau
CISO	Chief Information Security Officer
COBIT	Control Objectives for Information and Related Technologies
CSF	Cybersecurity Framework
CVaR	Cyber Value at Risk
DDOS	Distributed Denial of Service
DHS	Department of Homeland Security
FAIR	Factor Analysis of Information Risk
FDIC	Federal Deposit Insurance Corporation
FFIEC	Federal Financial Institutions Examination Council
FFRDC	Federally Funded Research and Development Center
FRB	Board of Governors of the Federal Reserve System
FS-ISAC	Financial Services Information Sharing and Analysis Center
FSOC	Financial Stability Oversight Council
FSS	Financial Services Sector
H&S	Hubbard and Seiersen
HSSEDI	Homeland Security Systems Engineering & Development Institute
IDART	Information Design Assurance Red Team
ISO	International Organization for Standardization
IT	Information Technology
NCUA	National Credit Union Administration

Acronym	Definition
NGCI	Next Generation Cyber Infrastructure
NIST	National Institute of Standards and Technology
OCC	Office of the Comptroller of the Currency
OCTAVE	Operationally Critical Threat, Asset, and Vulnerability Evaluation
OFR	Office of Financial Research
ROI	Return on Investment
SOX	Sarbanes–Oxley Act of 2002
S&T	Science and Technology Directorate
STIX	Structured Threat Information eXpression
STRIDE	Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege
SWIFT	Society for Worldwide Interbank Financial Telecommunication
VaR	Value at Risk

List of References

1. Armed Forces Communications and Electronics Association (AFCEA). 2013. "The Economics of Cybersecurity: A Practical Framework for Cybersecurity Investment," October 2013. <http://www.afcea.org/committees/cyber/documents/CyberEconfinal.pdf>
2. Armed Forces Communications and Electronics Association (AFCEA). 2014. "The Economics of Cybersecurity Part II: Extending the Cybersecurity Framework," April 2014. <http://www.afcea.org/committees/cyber/documents/EconomicsofCybersecurityPartII-Final4-2-14.pdf>
3. Basel Committee on Banking Supervision (BCBS). 1996a. Press Release on Market Risk, 10 December 1996. <https://www.bis.org/press/p961210.htm>
4. Basel Committee on Banking Supervision (BCBS). 1996b. "Overview of the Amendment to the Capital Accord to Incorporate Market Risks," January 1996. <https://www.bis.org/publ/bcbs23.pdf>
5. Beckstrom, R. 2014. "CyberVaR: Quantifying the risk of loss from cyber attacks," December 2014. <http://www.beckstrom.com/uncategorized/cybervar-quantifying-risk-loss-cyber-attacks/>
6. Black, P. E., Scarfone, K., and Souppaya, M. 2008. "Cyber security metrics and measures." In *Wiley Handbook of Science and Technology for Homeland Security*, J. G. Voeller, Ed. John Wiley & Sons, 2015.
7. Bodeau, D. and Graubart, R. 2014. "A Framework for Describing and Analyzing Cyber Strategies and Strategic Effects," MTR 140346, PR 14-3407. The MITRE Corporation, Bedford, MA. 2014.
8. Box, G. E. P. 1979. "Some Problems of Statistics and Everyday Life." *Journal of the American Statistical Association*, Vol. 74, No. 365, 1979.
9. Buith, J., and Spataru, D. 2015. "The Benefits, Limits of Cyber Value-at-Risk." CIO Insights and Analysis from Deloitte. *CIO Journal, The Wall Street Journal*, May 2015. <http://deloitte.wsj.com/cio/2015/05/04/the-benefits-limits-of-cyber-value-at-risk/>
10. Center for Internet Security (CIS). 2009. "The CIS Security Metrics." Consensus Metric Definitions, v1.0.0, May 11, 2009.
11. Dale, R. 1996. "Regulating the new financial markets." In *The Future of the Financial System*, M. Edey, Ed. Proceedings of a Conference, Reserve Bank of Australia, 1996, pp. 215-245. <https://www.rba.gov.au/publications/confs/1996/pdf/conf-vol-1996.pdf>
12. Duffie, D., and Pan, J. 1997. "An Overview of Value at Risk." *Journal of Derivatives*, Vol. 4, No. 3, Spring 1997.
13. FAIR Institute. <http://www.fairinstitute.org/>
14. Federal Financial Institutions Examination Council (FFIEC). Cybersecurity Assessment Tool. <https://www.ffiec.gov/cyberassessmenttool.htm>
15. Garbade, K. D. 1986. "Assessing risk and capital adequacy for Treasury securities." *Topics in Money and Securities Markets*, 22, New York: Bankers Trust.

16. Garbade, K. D. 1987. "Assessing and allocating interest rate risk for a multi-sector bond portfolio consolidated over multiple profit centers." *Topics in Money and Securities Markets*, 30, New York: Bankers Trust.
17. Geer, D., and Pareek, M. Index of Cyber Security.
<http://www.cybersecurityindex.org/index.php>
18. Geer, D., and Pareek, M. 2016. "The Index of Cyber Security-Detailed Report." Index of Cyber Security, October 2016.
19. Holton, G. A. 2004. "Defining risk." *Financial Analysts Journal*, 60(6), 19-25.
20. Holton, G. A. 2014. *Value-at-Risk: Theory and Practice*, Second Edition.
<https://www.value-at-risk.net/>
21. Hubbard, D. W., and Seiersen, R. 2016. *How to Measure Anything in Cybersecurity Risk*. John Wiley & Sons, 2016.
22. ISACA and Cybersecurity Nexus (CSX). 2015. "2015 Global Cybersecurity Status Report—US Data," January 2015. https://www.isaca.org/cyber/Documents/2015-US-Data-Sheet-for-Global-Cybersecurity-Status-Report_mkt_Eng_0115.pdf
23. Josey, A., et al. 2014. "An Introduction to the Open FAIR Body of Knowledge." The Open Group, June 25, 2014. <https://publications.opengroup.org/w148>
24. Markowitz, H. 1952. "Portfolio selection." *The Journal of Finance*, 7(1), 77-91.
25. Miller, J. 2016. "What is the value of a cyber investment?" Federal News Radio, July 25, 2016. <http://federalnewsradio.com/reporters-notebook-jason-miller/2016/07/value-cyber-investment/>
26. NIST. 2011. "Managing Information Security Risk: Organization, Mission, and Information System View," NIST Special Publication 800-39, March 2011.
<http://csrc.nist.gov/publications/nistpubs/800-39/SP800-39-final.pdf>
27. NIST. 2012. "Guide for Conducting Risk Assessments," NIST Special Publication 800-30 Rev.1, September 2012. <https://csrc.nist.gov/publications/detail/sp/800-30/rev-1/final>
28. NIST. 2014. "Framework for Improving Critical Infrastructure Cybersecurity," Version 1.0, February 12, 2014. <http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214.pdf>
29. Office of Financial Research (OFR). 2015. "Financial Stability Report," 2015.
https://financialresearch.gov/financial-stability-reports/files/OFR_2015-Financial-Stability-Report_12-15-2015.pdf
30. Object Management Group (OMG). 2014. "UML Operational Threat & Risk Model Request for Proposal," OMG Document SysA/2014-06-17. <https://www.omg.org/cgi-bin/doc.cgi?sypa/2014-6-17>
31. Payne, S. 2006. "A Guide to Security Metrics." SANS Institute, 2007.
<https://www.sans.org/reading-room/whitepapers/auditing/guide-security-metrics-55>
32. Ponemon Institute. 2016. "2016 Cost of Data Breach Study: United States." Ponemon Institute Research Report. IBM, June 2016. <http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=sel03094usen>

33. PricewaterhouseCoopers. 2015. "US cybersecurity: Progress stalled – Key findings from the 2015 US State of Cybercrime Survey," July 25, 2015. <http://www.pwc.com/us/en/increasing-it-effectiveness/publications/assets/2015-us-cybercrime-survey.pdf>
34. Reagan, J. R., Raghavan, A., and Thomas, A. 2016. "Quantifying risk: What can cyber risk management learn from the financial services industry?" Deloitte Review, issue 19, July 25, 2016. <http://dupress.com/articles/quantifying-risk-lessons-from-financial-services-industry>
35. RiskLens. <http://www.risklens.com/>
36. Sandia National Laboratories. Information Design Assurance Red Team. <http://idart.sandia.gov/>
37. World Economic Forum (WEF). 2015. "Partnering for Cyber Resilience: Towards the Quantification of Cyber Threats," January 2015. http://www3.weforum.org/docs/WEFUSA_QuantificationofCyberThreats_Report2015.pdf