Cyber100 Compass: Quantification of Cybersecurity Risks for Systems Transitioning to High Levels of Renewables (Final Report)

Maurice Martin, Chelsea Neely, Rebecca Hanes, and Laura Leddy

National Renewable Energy Laboratory

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5R00-88096
July 2024

Cyber100 Compass: Quantification of Cybersecurity Risks for Systems Transitioning to High Levels of Renewables (Final Report)

Maurice Martin, Chelsea Neely, Rebecca Hanes, and Laura Leddy

National Renewable Energy Laboratory

Suggested Citation

Martin, Maurice, Chelsea Neely, Rebecca Hanes, and Laura Leddy. 2024. *Cyber100 Compass: Quantification of Cybersecurity Risks for Systems Transitioning to High Levels of Renewables (Final Report)*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5R00-88096. https://www.nrel.gov/docs/fy24osti/88096.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5R00-88096
July 2024

National Renewable Energy Laboratory
15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Electricity and Office of Cybersecurity, Energy Security, and Emergency Response. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Acknowledgments

The authors gratefully acknowledge the U.S. Department of Energy Office of Electricity (OE) and Office of Cybersecurity, Energy Security, and Emergency Response (CESER) for the funding support and feedback provided throughout the project. In addition, the authors thank Greg Wyss of Sandia National Laboratories for his review as well as the following National Renewable Energy Laboratory staff for their reviews: Tim Coleman, Dylan Hettinger, Jordan Henry, and Tami Reynolds. Finally, we thank the subject matter experts from the following organizations who generously provided data, insights, and encouragement in this effort: the American Council on Renewable Energy, the American Public Power Administration, the Black & Veatch Corporation, the CIP Corps LLC, the National Rural Electric Cooperative Association, the Platte River Power Authority, and Southern California Edison.

List of Acronyms

CAIDI	Customer Average Interruption Duration Index
CDF	Customer Damage Function
CECA	Clean Energy States Alliance
DER	distributed energy resource
DER-CF	Distributed Energy Resource Cybersecurity Framework
DHS	U.S. Department of Homeland Security
DOE	U.S. Department of Energy
EIA	U.S. Energy Information Administration
ICE	Interruption Cost Estimate
IT	information technology
NERC	North American Electric Reliability Corporation
NIST	National Institute of Standards and Technology
NREL	National Renewable Energy Laboratory
OSHA	Occupational Safety and Health Administration
OT	operational technology
ReEDS	Regional Energy Deployment System
REL	recommended expectation of loss
SAIDI	System Average Interruption Duration Index
SAIFI	System Average Interruption Frequency Index
SME	subject matter expert

Executive Summary

The shift to high deployment levels of renewable generation will entail a significant reengineering of the electric grid. Future grids will incorporate more distributed energy resources, which will likely be managed by sophisticated control algorithms operating over communication networks reaching almost (or perhaps all the way) to the grid edge to maintain stability and reliability. Also, the many different generation facilities, storage facilities, and responsive loads will be managed by multiple operating entities.

Transitioning to high levels of renewables requires significant investment. As investors, utilities, customers, and others prepare for the clean energy transition, there is a need to understand how restructuring the grid to accommodate renewables will change the attack surface of the grid and the accompanying cyber risks. Today, however, the cyber-physical risks associated with electric grids incorporating high deployment levels of renewables remain largely unknown.

The Cyber100 Compass proof-of-concept application attempts to quantify future cyber-physical security risks by combining risk data gathered from subject matter experts (SMEs) with input from system planners about conditions they expect to be true about their electric systems in the future.

In creating the Cyber100 Compass, researchers developed tools and techniques for collecting input from both the application users (system planners) and the SMEs. Users provide data about their organization’s tolerance for risk, the value they place on avoiding the consequences of different cyber events, and the conditions that they expect to be true on their systems at some point in the future. The SMEs provide baseline probabilities for different cyber events; the probability that an event will be low, moderate, or high impact; and the amount by which user-identified conditions on their systems will change the likelihood of these cyber events. The application takes both the user and SME inputs and performs a series of Monte Carlo simulations to arrive at a quantification of cybersecurity risk.

The application in its current form is not sufficiently mature for use by utilities. Several known issues and open questions need to be addressed to increase its usability. Future research to improve the framework could include creating cost-effective ways to increase SME engagement and input; addressing the limits of some of the employed mathematical methods; improving techniques for understanding and addressing consistency in the SME input; developing methods for validating results; and identifying and integrating other types of risk data, such as historical data or data generated through experimentation. Some of the challenges discovered in the creation of the Cyber100 Compass framework are intertwined with larger challenges regarding cybersecurity risk analysis in situations where data are sparse.

Despite these challenges, the Cyber100 Compass’s proof-of-concept application represents an important first step toward developing much-needed guidance for utility planners about the cyber risks they face as they transition their electric grids to high levels of renewables. The Cyber100 Compass offers a promising and novel approach to quantifying risks for future energy systems—an understudied, poorly understood, and increasingly critical area of cyber risk management.

Table of Contents

List of Figures

List of Tables

1 Introduction

As of 2022, renewable energy accounts for approximately 21% of the total utility-scale electricity generation in the United States (EIA 2023). The prices of renewable generation technologies have continued to decrease, and a wide range of jurisdictions have established some sort of renewable portfolio standards or clean energy standards. Currently, 22 states plus the District of Columbia and Puerto Rico have 100% clean energy goals, with target dates ranging from 2030 to 2050 (CESA 2023). Policy and market forces will likely continue to drive the increased integration of renewable energy into our electric systems.

The shift to high deployment levels of renewables will entail a significant reengineering of the electric grid. Future renewable grids will incorporate more distributed energy resources (DERs), such as wind turbines, solar panels, and battery energy storage systems. These will likely be managed by sophisticated control algorithms operating over communication networks reaching almost (or perhaps all the way) to the grid edge to maintain stability and reliability. Also, the many different generation facilities, storage facilities, and responsive loads could be managed by multiple operating entities. In addition to the utility, these operating entities might include owners of large-scale solar, wind, or storge facilities; aggregators of many small-scale generation units, such as rooftop solar; building management systems; operators of electric vehicle charging stations; etc. Although the current grid is sometimes discussed as a system of systems, the new renewable grid will embody this idea to a much greater degree and on a much larger scale. The individual operating entities will be the systems in this scenario, each with their own operational infrastructures, control loops, and business objectives. The system of systems describes the collective behavior of the grid as these constituent systems are brought together in a purposeful, coordinated way.

At the same time that these changes are being planned and implemented, the threat landscape is also rapidly evolving. Cyberattacks that produce physical consequences are becoming less rare—Stuxnet (Zetter 2014) and the power outages on the Ukrainian grid (Greenberg 2017) are frequently cited examples. But the NotPetya attack also produced physical disruptions that included $200 million in losses to shipping company A.P. Moller-Maersk (Mathews 2017) due to the inoperability of ports and other facilities. The 2021 ransomware attack on the Colonial Pipeline shows how cyberattacks can disrupt physical processes even when the attacks do not reach the operational technology (OT) systems—Colonial shut down its OT systems to prevent the malware from spreading to them, resulting in the disruption of retail gas supply chains (Wood 2023). Although no incidents have yet been cited, a recently discovered malware, Pipedream, offers adversaries an advanced tool kit that can disrupt or take control of almost any industrial control systems environment (Greenberg 2022). Cyberattacks with physical consequences are now part of the threat landscape and must be factored into risk assessments and budgetary allocations for critical infrastructure.

Transitioning to high levels of renewables requires significant investment. As investors, utilities, customers, and others prepare for clean energy transitions, there is a need to understand how restructuring the grid to accommodate renewables will change the attack surface of the grid and the accompanying cyber risks. Today, however, the cyber-physical risks associated with electric grids incorporating high deployment levels of renewables remain largely unknown. As deployments increase, utilities and other system planners could unintentionally build in systemic cyber vulnerabilities that would be difficult to retroactively address. The principle of *security by design*, an approach to software and hardware development that seeks to design systems as securely as possible from the start, is widely accepted at the device level and must also be applied at the level of a system of systems.

Traditional risk matrix approaches measure cybersecurity risks using ordinal scales of measurement—for example, red, yellow, and green notations (sometimes called the "stoplight" approach) might represent low, moderate, and high risk based on an event's likelihood and impact. Although these qualitative assessments are widely used, it is unclear whether these approaches help reduce risks or are more effective than expert intuition alone (Hubbard and Seiersen 2016, 85). Some studies conclude that because qualitative assessments are subjective and open to interpretation, they might have a worse-than-random impact on decision makers (Cox 2008).

The Cyber100 Compass proof-of-concept application uses an alternate methodology to quantify future cyber-physical security risks by eliciting risk data from subject matter experts (SMEs)—those who have expertise in the fields of power systems, cybersecurity, and risk management. SME inputs are captured in data tables that drive risk quantification in the application, allowing for reuse across systems of many sizes, configurations, and generation mixes. The application quantifies risks in terms of monetary values corresponding to losses from different types of cyber events within a typical 12-month period. Although the application is not currently mature enough for recommended use by utilities, it demonstrates novel data gathering and mathematical approaches that could be used to assess the impacts of cyberattacks on future power systems with high deployment levels of renewables.

1.1 Cyber100 Compass Framework Design Objectives

The Cyber100 Compass framework was created to help application users—system planners, operators, and other stakeholders—improve their understanding of the cyber risks they face as they transition to high levels of renewable energy. To achieve the main objectives of the Cyber100 Compass, the framework and proof-of-concept application were designed to account for the following scopes and constraints:

1. **Questions must be answerable by the user.** The Cyber100 Compass should only request information that the users can reasonably be expected to provide. For example, asking users "Which techniques will you employ against quantum-based attacks?" is not useful because system planners and operators have little information regarding which solutions will be available to them once quantum-based attacks are a concern. In contrast, questions about system architectures should be answerable because such information should be part of system upgrade plans.
2. **It must be possible to generate the data necessary for back-end calculations.** The application must be able to determine how grid conditions impact risk. The data quantifying these impacts for back-end calculations were created based on elicitations from SMEs. A data gathering tool allowed the SMEs to review and verify the impacts of the data they provided. In addition, the data must be structured for easy integration into the Cyber100 Compass data tables.

3. **The output of the framework should express risks quantitatively rather than qualitatively.** Qualitative outputs (e.g., red, yellow, and green indicators) have been shown to have limited value and might even have a worse-than-random impact on decision makers (Cox 2008). The monetary expression of risks is more effective because it allows cyber risks to be integrated into organization-wide risk management efforts and appropriately budgeted.
4. **The framework should be extensible in terms of the conditions input by users.** As explained in Section 3.3, users provide answers to questions about conditions they expect to be true about their systems as they integrate high levels of renewable energy. The framework should allow for the addition of new condition questions over time as new technologies, cybersecurity risks, or other concerns arise.
5. **The framework should focus on attacks on OT systems that produce physical effects.** Tools already exist to assess cyber risks to information technology (IT) systems. Cyber risks for OT systems, which can produce physical effects, are less studied. Table 1 lists the physical effects that the Cyber100 Compass focuses on along with examples.

Table 1. Possible Physical Effects of Cyberattacks

Physical Effect	Example
Power outage	An attacker gains control of substation equipment and opens circuit breakers, shutting off power.
Harm to equipment	An attacker overloads a transformer, causing it to overheat and fail.
Harm to employees (of the utility or operating entity)	An attacker energizes a line while it is under repair, exposing workers to electric shock.
Harm to community	An attacker sends an excessive amount of current over a conductor, annealing it and starting a wildfire.
Loss of productivity or efficiency	A denial-of-service attack on a utility communication system prevents DERs from operating at optimal efficiency.

The Cyber100 Compass project team built a proof-of-concept application that instantiates the Cyber100 Compass framework. Through coding the application, the team thought through details of operationalizing this new risk assessment approach and improved the methodology and usability of the framework. Section 9 explains some of the challenges the project team faced in developing the framework and the limitations of the application in its current form.

2 Conceptual Foundation

2.1 Sparse Data Analytics

The Cyber100 Compass analysis is an example of *sparse data analytics*, meaning that the analysis is done using relatively small amounts of data.

Historical data for cyber risks are sparse in general (Sheehan et al. 2021) because cyberattacks are a relatively new phenomena. (In contrast, for example, the insurance industry can draw on hundreds of years of weather data when setting rates for flood insurance.) Historical data are especially sparse for the specific problem of assessing future risks from the transition to high levels or renewables. Where historical data exist, they might be of limited value because cyber threats, cyber vulnerabilities, and cyber defenses change quickly (CrowdStrike 2023). Last, historical data might only include data about attacks on IT systems, leaving out OT attacks that cause physical impacts relevant to this framework. Historical data were considered but ultimately were not used for the Cyber100 Compass project due to these limitations; however, as more historical data become available over time, follow-on work could enable the Cyber100 Compass risk assessment framework to incorporate them.

The conceptual groundwork for quantitative assessments of cyber risks using sparse data from other sources was explored in the book *How to Measure Anything in Cybersecurity Risk* by Douglas W. Hubbard and Richard Seiersen. This information proved highly instructive when developing the Cyber100 Compass framework. Hubbard and Seiersen define the measurement of cyber risks as the process of making observations that reduce uncertainty about cyber risk and expressing those observations as data (Hubbard and Seiersen 2016, 24). Confronted with the problem of sparse data, they observe that:

- "You probably need less data than your intuition tells you…"
- You often have more data than you think you have.
- You can update your predictive ability as you get more data (Hubbard and Seiersen 2016, 34, 59, 201).

Advancing cyber risk analysis is therefore the process of finding which data are available and fitting them into an appropriate analytic framework—one that allows for the addition of more data as they become available.

2.2 Focus on Operational Technology

The Cyber100 Compass focuses on cyber events affecting systems and networks that control physical devices and processes—for example, a malicious command on an electric utility OT network might open a circuit breaker and cause an outage. (See Table 1.)

Attacks focused on IT systems are out of scope for the Cyber100 Compass. IT systems store, transmit, and process information, but they do not control physical devices. An example of an IT system is one designed to process financial records. In some cases, attacks might begin on IT systems and pivot to OT systems. In these scenarios, the Cyber100 Compass is concerned with only the OT impacts of such attacks.

2.3 Front-End Data Versus Back-End Data

The Cyber100 Compass method for quantifying risks depends on decomposing the risks into several elements, then performing Monte Carlo simulations to calculate losses. The decomposition of the risk results in two distinct sets of data: front-end data supplied by the application users and back-end data supplied by qualified SMEs. The elements of these data sets are explained in detail in Section 3 and Section 4.

3 Front-End Data

Section 3 describes the Cyber100 Compass data gathering concepts and methods that users encounter on the front end, or user interface, of the proof-of-concept application. The Cyber100 Compass elicits three categories of data from users: risk tolerance data, event avoidance value data, and condition selection data.

3.1 Risk Tolerance Data

Risk tolerance describes an organization’s willingness to accept certain levels of risk based on the financial losses that could occur from cyberattacks in a typical year. Risk tolerance inputs in the Cyber100 Compass help users quantify the potential losses their organization might be willing to accept from certain kinds of cyberattacks.

Risk tolerance combines acceptable losses from cyber events with the probability of that level of loss occurring. Risk tolerance values, compared to simulated losses derived from other front-end data components, are used to generate a risk tolerance curve to show users whether their system development plans are leading them toward unacceptable levels of risk.

The risk tolerance curve, a concept widely used in decision sciences (Hubbard and Seiersen 2016, 47) visualizes acceptable losses based on the probability of occurrence, with larger losses less likely than smaller losses. Figure 1 shows an example of a risk tolerance curve, with annotations noting two illustrative points on the curve. As indicated, this sample curve shows that an organization is willing to accept an approximately 35% probability of an annual loss of $10,000 or more and an approximately 15% probability of an annual loss of $100,000 or more.

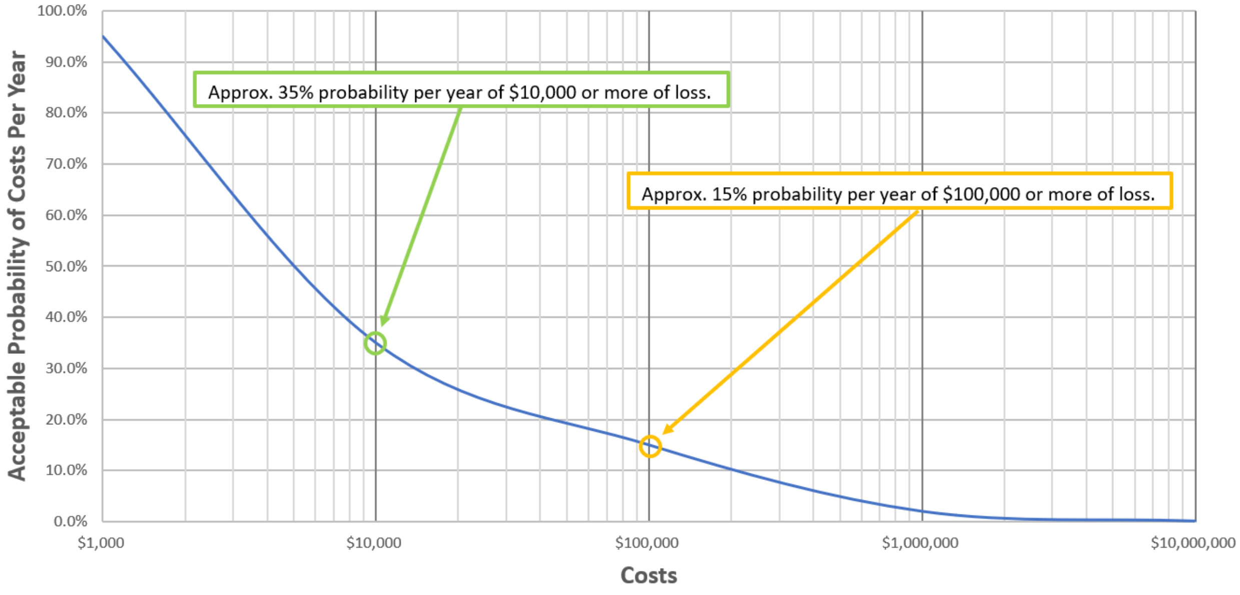

Figure 1. Example risk tolerance curve

3.2 Event Avoidance Value Data

Event avoidance values allow users to quantify the value they place on avoiding cyber events addressed in the Cyber100 Compass. In the resilience space, these values are often called "avoided costs" (NREL 2022). The proof-of-concept application considers five types of cyber events that create physical effects. (See Table 1.)

Avoided costs are the estimated costs that would result from a possible cyberattack, which users hope to avoid. For instance, if a utility knows that an outage of its entire system lasting 12 hours would result in financial damages of $100,000, the utility can act to prevent the outage, and those actions would result in an avoided cost of $100,000.

When calculating the avoided cost, the Cyber100 Compass asks users to include all possible costs that might arise from the attack. At a minimum, these could include the loss of revenue from business that cannot be transacted due to the attack as well as the cost of recovery efforts, but there are many other possible costs to consider. For instance, utilities might want to include the economic impact on their customers of a cyberattack-induced outage. The Cyber100 Compass allows for considerable flexibility in calculating the avoided costs for different cyber events based on the unique characteristics and circumstances of the electric system being assessed. The tool suggests some cost factors that users can consider, but, ultimately, deciding which costs to include is up to the user.

For some types of events, users might partially base valuations on similar past events with a non-cyber origin. For instance, when placing a value on avoiding an outage, a utility can look at data from outages caused by weather, equipment failure, etc. The utility does not need to have experienced an outage due to a cyberattack—much of the value is determined by the nature of the event, not the cause itself; however, the cause of the attack might influence how the utility estimates the cost of recovery.

Each event is divided into three impact levels: low, moderate, and high. The impact levels represent a simplified assessment scale inspired by the National Institute of Standards and Technology (NIST) Special Publication 800-30, *Guide to Conducting Risk Assessments*, Appendix H, Table H-3 (NIST 2012b). The Cyber100 Compass leveraged this assessment scale to guide users in estimating the avoidance costs they must input across the five event categories. Table 2 shows an example of an event input—in this case, in the category of harm to equipment—that users are asked to complete.

Table 2. User Interface Input for the “Harm to Equipment” Event Avoidance Values

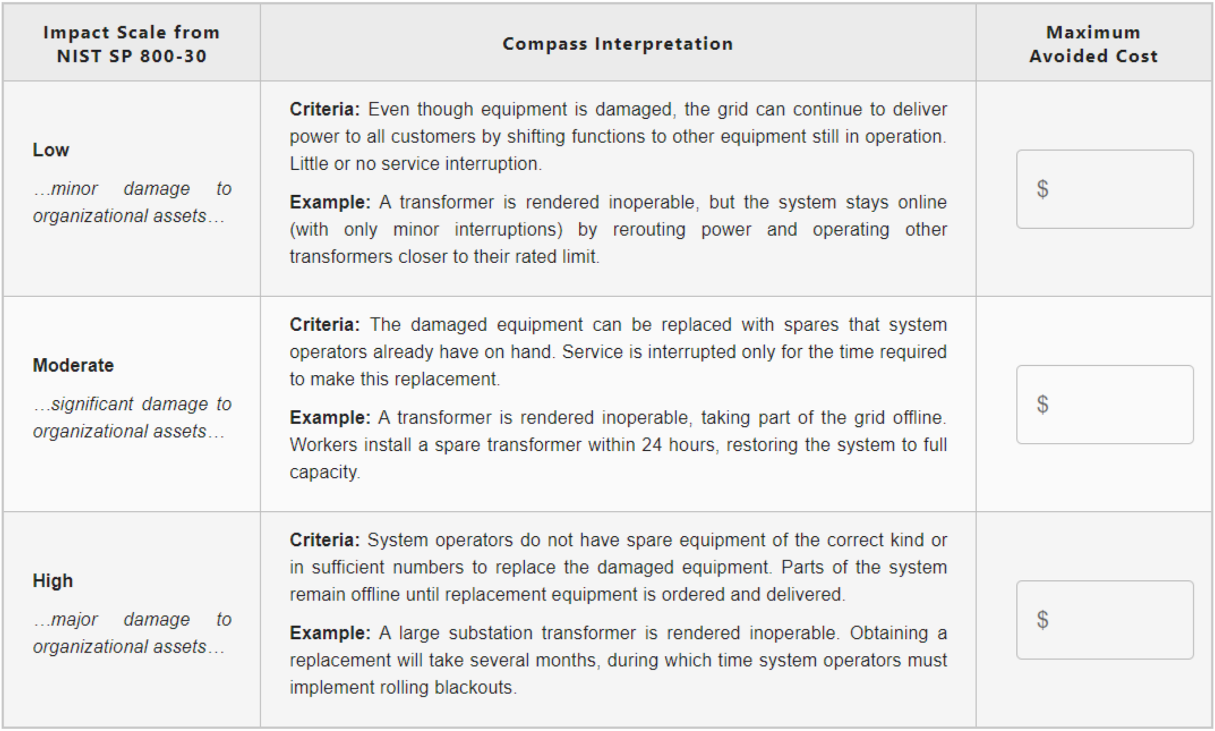

Impact Scale from NIST SP 800-30	Compass Interpretation	Maximum Avoided Cost
Low *…minor damage to organizational assets…*	**Criteria:** Even though equipment is damaged, the grid can continue to deliver power to all customers by shifting functions to other equipment still in operation. Little or no service interruption. **Example:** A transformer is rendered inoperable, but the system stays online (with only minor interruptions) by rerouting power and operating other transformers closer to their rated limit.	$
Moderate *…significant damage to organizational assets…*	**Criteria:** The damaged equipment can be replaced with spares that system operators already have on hand. Service is interrupted only for the time required to make this replacement. **Example:** A transformer is rendered inoperable, taking part of the grid offline. Workers install a spare transformer within 24 hours, restoring the system to full capacity.	$
High *…major damage to organizational assets…*	**Criteria:** System operators do not have spare equipment of the correct kind or in sufficient numbers to replace the damaged equipment. Parts of the system remain offline until replacement equipment is ordered and delivered. **Example:** A large substation transformer is rendered inoperable. Obtaining a replacement will take several months, during which time system operators must implement rolling blackouts.	$

Note that the user interface specifically asks for the “maximum avoided cost.” The maximum avoided cost describes the highest dollar amount a user estimates the organization could lose at each impact level for each type of event. Each of the five event categories requires users to input the maximum avoided costs at low-, moderate-, and high-impact levels. Once users complete all the inputs across all five event categories, they can proceed to the next part of the tool—conditions.

3.3 Condition Selection Data

The Cyber100 Compass asks users to select the conditions that describe any constraints, resources, requirements, controls, or other factors that modify the cybersecurity risks of a system’s energy transformation plans. Conditions are presented as a series of questions about the utility’s energy transition plans, most often looking 5 years into the future. These questions allow users to describe aspects of their current and future energy systems that will impact cybersecurity risks.

The conditions in the Cyber100 Compass are similar to NIST’s definition of a “predisposing condition” (NIST n.d.):

> A condition that exists within an organization, a mission/business process, enterprise architecture, or information system including its environment of operation, which contributes to (i.e., increases or decreases) the likelihood that one or more threat events, once initiated, will result in undesirable consequences or adverse impact to organizational operations and assets, individuals, other organizations, or the Nation.

The Cyber100 Compass extends this idea from the present into the future. Many of the condition questions ask users what they expect to be true about their electric system within the next 5 years.

These questions are meant to elicit information; they are not meant to be—and they should not be interpreted as—policy recommendations. Further, the Cyber100 Compass is not primarily focused on the physical controls for security. The purpose of the Cyber100 Compass is to assess the risks of cyberattacks producing physical impacts on energy systems, not physical security breaches that have cybersecurity impacts.

In the proof-of-concept application, condition questions are spread across the following five categories:

- Changes to grid topology
- Changes to system-of-systems architecture
- Communications
- Security controls
- Regulatory environment.

Each answered condition question refines the user's risk estimate specific to the system under consideration. These refinements make the probabilistic calculations that quantify the cybersecurity risks that are more specific to the assessed future energy system.

Because the Cyber100 Compass users are asked about future aspects of their grids, inevitably, there is some amount of uncertainty. Users of the Cyber100 Compass are not expected to have perfect knowledge regarding their energy systems within the 5-year time span identified in most questions. If a question asks for only one answer, the Cyber100 Compass gives users the opportunity to select "unsure"; however, users are encouraged to use this option as infrequently as possible. Selecting more decisive answers—based on business plans, information from third parties, trends, the user's own foresight, and other sources—produces results that are more specific to their future system. The more questions that are decisively answered, the better the Cyber100 Compass can estimate risks and calculate losses from cybersecurity events.

Though not exhaustive, the following resources might be useful to users when answering conditions questions:

- Utility personnel, which could include the chief security officer, the chief technology officer, the chief information officer, system planners, control engineers, power engineering teams, communication engineers, the chief financial officer, the corporate governance team, the government affairs office, etc.
- System planning data, which could include design documents, power purchase agreements, capacity expansion models, projections of load growth, etc.
- Persons involved in the development and implementation of interconnection agreements
- Current or future third-party operating entities
- State public utility commissions
- State, regional, or national service organizations, for example, the National Rural Electric Cooperative Association or one of its state-level associations.

The project invested considerable time in designing conditions clearly. Each condition input includes the following headings:

- **Background:** Provides the necessary context for the question
- **Assumption(s):** Gives a qualitative explanation of how the input will affect the risk estimate
- **Question:** The actual request for information
- **References:** Provides citations and additional background information when needed to help the user understand the context of the question. (Note: Not all questions include references.)

Note that the project team's assumptions about how a condition input would affect the risk estimate did not always align with the data gathered from the SMEs. These headings would need to be revised if the proof-of-concept were to be developed further.

Appendix E provides the text of all the condition inputs. Table 3 shows an example condition question a user would answer. The response choices were generally arranged as a progression, with those likely to increase risks at the top and those likely to reduce risks at the bottom.

Table 3. Cyber100 Compass Condition Input Example: Degrees of Centrality

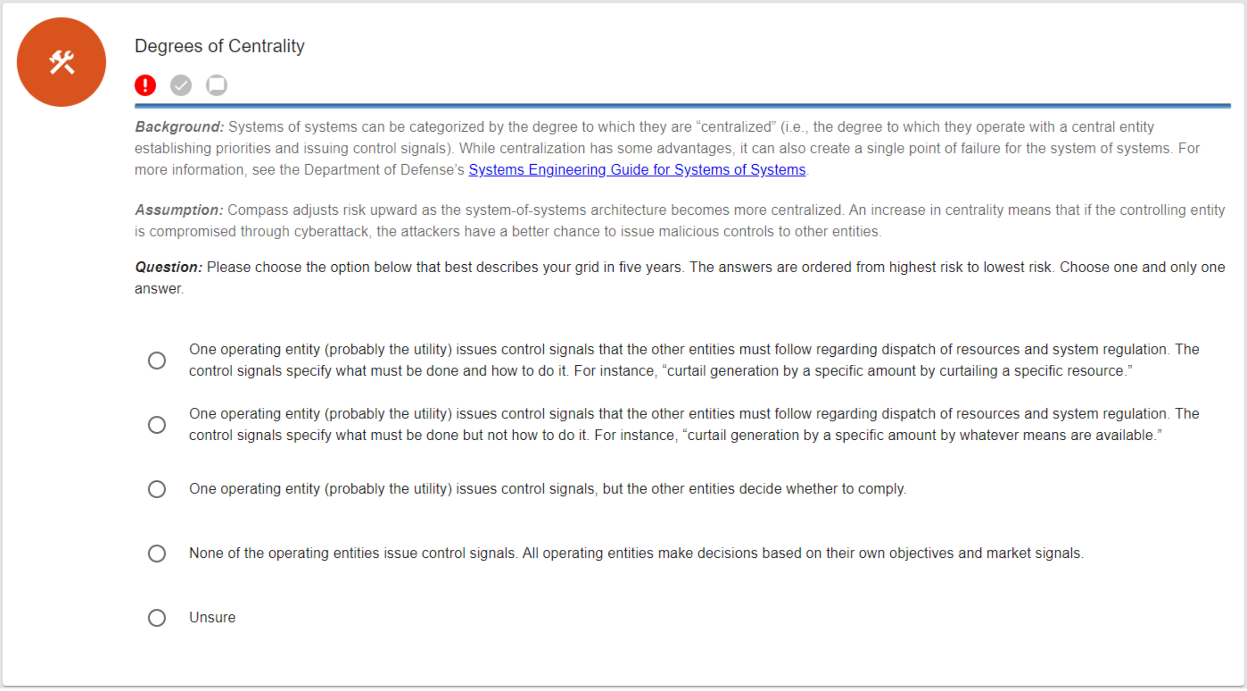

Degrees of Centrality

Background: Systems of systems can be categorized by the degree to which they are "centralized" (i.e., the degree to which they operate with a central entity establishing priorities and issuing control signals). While centralization has some advantages, it can also create a single point of failure for the system of systems. For more information, see the Department of Defense's Systems Engineering Guide for Systems of Systems.

Assumption: Compass adjusts risk upward as the system-of-systems architecture becomes more centralized. An increase in centrality means that if the controlling entity is compromised through cyberattack, the attackers have a better chance to issue malicious controls to other entities.

Question: Please choose the option below that best describes your grid in five years. The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

○ One operating entity (probably the utility) issues control signals that the other entities must follow regarding dispatch of resources and system regulation. The control signals specify what must be done and how to do it. For instance, "curtail generation by a specific amount by curtailing a specific resource."

○ One operating entity (probably the utility) issues control signals that the other entities must follow regarding dispatch of resources and system regulation. The control signals specify what must be done but not how to do it. For instance, "curtail generation by a specific amount by whatever means are available."

○ One operating entity (probably the utility) issues control signals, but the other entities decide whether to comply.

○ None of the operating entities issue control signals. All operating entities make decisions based on their own objectives and market signals.

○ Unsure

4 Back-End Data

The risk tables that comprise the back-end data include three components: baseline probabilities, conditional probability adjustment factors, and distribution of impact. Back-end data for the proof-of-concept application was elicited from SMEs.

The National Renewable Energy Laboratory (NREL) engaged with a variety of SMEs—those who have expertise in the fields of power systems, cybersecurity, and risk management—to contribute to and create data tables for computing the back-end probabilities to make the Cyber100 Compass run. These experts also contributed feedback to improve the proof-of-concept application. A complete description of SME feedback is included in Appendix B.

These individuals were invited to a series of facilitated discussions and working sessions where they were briefed on the Cyber100 Compass and guided through an exercise to complete an NREL-created spreadsheet to capture baseline probabilities, conditional probability adjustment factors, and distribution of impact.

4.1 Baseline Probabilities

Baseline probabilities can be thought of as recommendations from the SMEs to the application users regarding the level of risk a hypothetical "average" utility should prepare for regarding the five types of cyberattacks addressed in the Cyber100 Compass. Although the expression of this "recommendation" looks very different in the tables used to gather SME data, the SME recommendation (using power outages as an example) can be paraphrased as:

> Knowing only that your system operates in the United States, I recommend that you prepare for an X% probability of a power outage due to cyberattacks occurring each year.

Each of the five cyberattacks on which the Cyber100 Compass focuses has its own baseline probability.

The baseline probability is a recommendation, not a prediction. This is because none of the SMEs who engaged with this project were comfortable making a prediction about the likelihood of occurrence of any particular cyber event (such as a cyber-induced power outage); however, SMEs were willing to make recommendations regarding how prepared a utility should be for different cyber events. The distinction between a recommendation and a prediction is important because the baseline probability is the starting point for all calculations performed on the Cyber100 Compass back end. The conditions that the user selects serve to modify the baseline probability either up or down, reflecting the impact of the conditional risk—essentially modifying the original recommendation. The combination of these modifications serves as the basis for a series of Monte Carlo simulations that produce a monetary value, which can also be considered a recommendation of type "It is recommended that you prepare your system for annual losses due to cyberattack-induced power outages of $X per year."

To elicit the baseline probability values from the SMEs, the baseline probability itself was decomposed into three factors. These factors are based on the NIST Special Publication 800-30 (NIST 2012). For each type of cyber event, the SMEs selected a rating for three factors:

- **C = adversary capability.** For a particular event (e.g., outage), what is the cyberattacker's level of expertise, resources, and opportunities to support the attack (Table D-3)?
- **T = adversary targeting.** For a particular event, to what degree is the cyberattacker specifically interested in disrupting the electric system or (more specifically) a particular system operator (Table D-5)?
- **V = vulnerability severity.** How vulnerable is the electric system operator to cyberattack (Table F-2)?

The qualitative rating (very low, low, moderate, etc.) for each factor corresponds to an assigned value ranging from 0.01 for very low to .99 for very high. The three assigned values are multiplied to obtain the baseline probability. The assigned values are mostly evenly spaced—for instance, low is 0.25, moderate is 0.5, and the option between them is 0.375; however, very low and very high were offset from 0 and 1 by 0.01 to avoid producing baseline probabilities of exactly 0 or 1.

Table 4 shows the spreadsheet inputs for the baseline probabilities. Note: This spreadsheet is a separate data collection tool for SMEs and is not part of the proof-of-concept application.

Table 4. SME Inputs for Baseline Probabilities

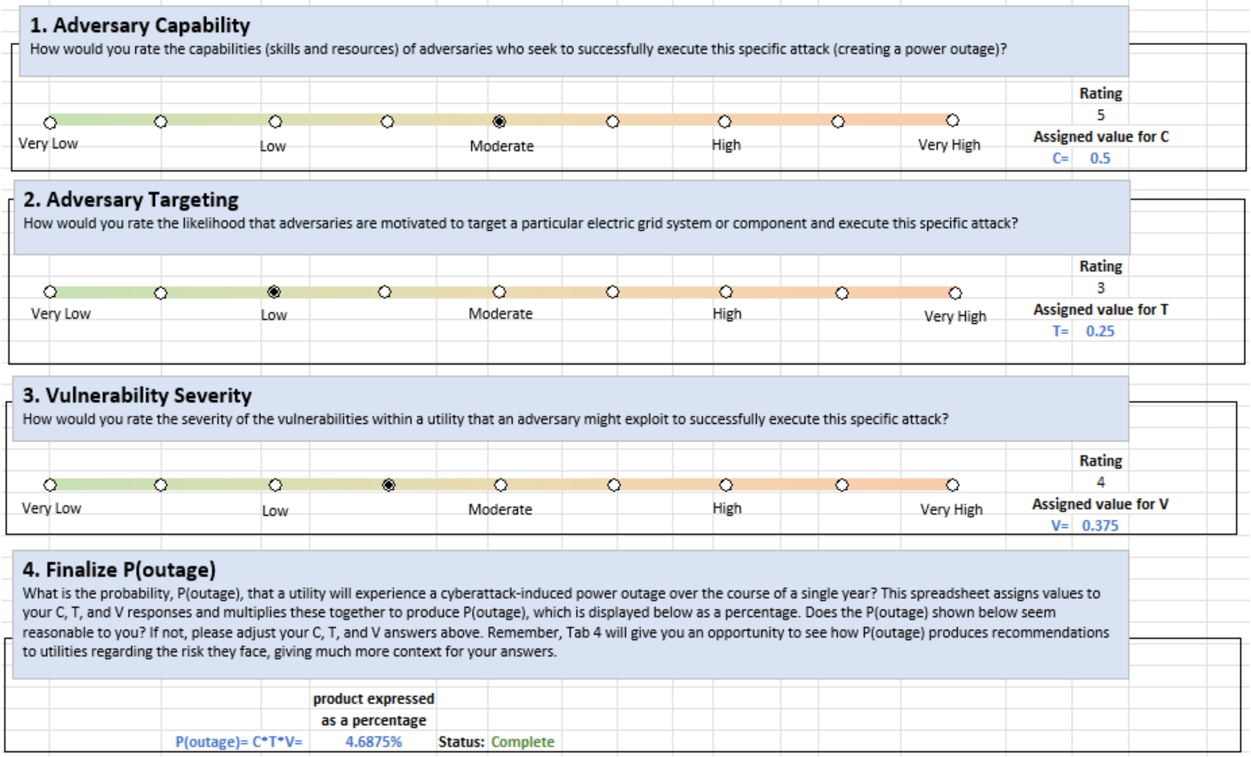

1. Adversary Capability

How would you rate the capabilities (skills and resources) of adversaries who seek to successfully execute this specific attack (creating a power outage)?

○ Very Low ○ ○ Low ○ ● Moderate ○ ○ High ○ ○ Very High

Rating
5
Assigned value for C
C= 0.5

2. Adversary Targeting

How would you rate the likelihood that adversaries are motivated to target a particular electric grid system or component and execute this specific attack?

○ Very Low ○ ● Low ○ ○ Moderate ○ ○ High ○ ○ Very High

Rating
3
Assigned value for T
T= 0.25

3. Vulnerability Severity

How would you rate the severity of the vulnerabilities within a utility that an adversary might exploit to successfully execute this specific attack?

○ Very Low ○ ○ Low ● ○ Moderate ○ ○ High ○ ○ Very High

Rating
4
Assigned value for V
V= 0.375

4. Finalize P(outage)

What is the probability, P(outage), that a utility will experience a cyberattack-induced power outage over the course of a single year? This spreadsheet assigns values to your C, T, and V responses and multiplies these together to produce P(outage), which is displayed below as a percentage. Does the P(outage) shown below seem reasonable to you? If not, please adjust your C, T, and V answers above. Remember, Tab 4 will give you an opportunity to see how P(outage) produces recommendations to utilities regarding the risk they face, giving much more context for your answers.

	product expressed as a percentage	
P(outage)= C*T*V=	4.6875%	Status: Complete

4.2 Distribution of Impact

The distribution of impact values assign probabilities that a single cyber event will be low, moderate, or high impact according to the descriptions of those levels provided to the users on the Event Avoidance Values input screens. The three impact probabilities—low, moderate, and high—must sum to one because the event (once determined to occur) must fall into one of those categories.

Each impact level for an event corresponds to a different range of loss values for that event. The Cyber100 Compass uses the distribution of impact probabilities during Monte Carlo simulations to determine which loss range to use to simulate the loss incurred from that event occurrence. (See Section 7 for details.)

Table 5 shows the spreadsheet inputs for the distribution of impact probabilities.

Table 5. SME Inputs for Distribution of Impact

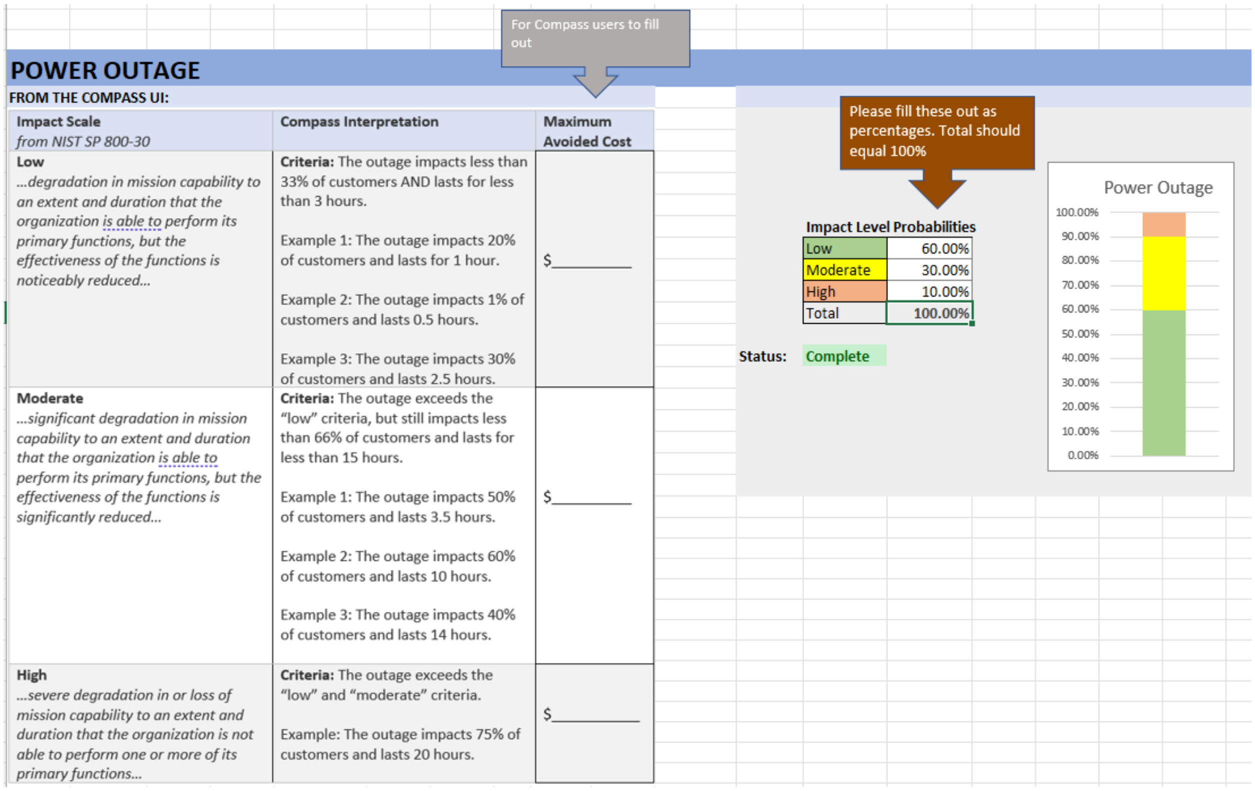

For Compass users to fill out

POWER OUTAGE

FROM THE COMPASS UI:

Impact Scale *from NIST SP 800-30*	Compass Interpretation	Maximum Avoided Cost
Low *...degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is noticeably reduced...*	**Criteria:** The outage impacts less than 33% of customers AND lasts for less than 3 hours. Example 1: The outage impacts 20% of customers and lasts for 1 hour. Example 2: The outage impacts 1% of customers and lasts 0.5 hours. Example 3: The outage impacts 30% of customers and lasts 2.5 hours.	$_________
Moderate *...significant degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is significantly reduced...*	**Criteria:** The outage exceeds the "low" criteria, but still impacts less than 66% of customers and lasts for less than 15 hours. Example 1: The outage impacts 50% of customers and lasts 3.5 hours. Example 2: The outage impacts 60% of customers and lasts 10 hours. Example 3: The outage impacts 40% of customers and lasts 14 hours.	$_________
High *...severe degradation in or loss of mission capability to an extent and duration that the organization is not able to perform one or more of its primary functions...*	**Criteria:** The outage exceeds the "low" and "moderate" criteria. Example: The outage impacts 75% of customers and lasts 20 hours.	$_________

Please fill these out as percentages. Total should equal 100%

Impact Level Probabilities

Low	60.00%
Moderate	30.00%
High	10.00%
Total	100.00%

Status: Complete

After the SMEs have provided baseline probabilities and distribution of impact values, the spreadsheet allows them to see how the values they provided translate into dollars. These dollar values are labeled REL for "recommended expectation of loss," defined as the amount that the SME recommends the user be prepared to lose due to a particular cyber event per year. The spreadsheet calculates single REL values for each event and separate values for each impact level.

Note that these calculations also require event avoidance values; however, because there is not actual user data at this point in the process, the spreadsheet used to capture the SME inputs includes sample event avoidance values based on a fictional utility. (See Appendix D for an example fictional utility.) SMEs can use the sample values as is or alter them and see the impact on the RELs in the resulting graphs. Table 6 illustrates this process.

Table 6. SME Inputs for Approving Preliminary Results Including Example Values

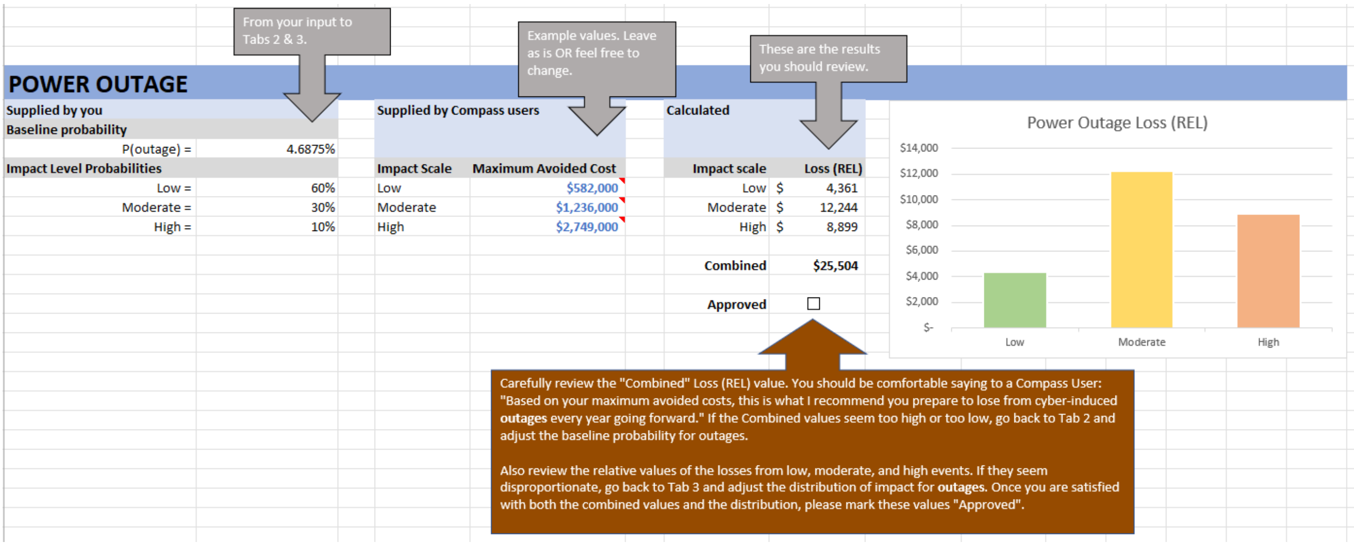

From your input to Tabs 2 & 3.

Example values. Leave as is OR feel free to change.

These are the results you should review.

POWER OUTAGE

Supplied by you		Supplied by Compass users		Calculated	
Baseline probability					
P(outage) =	4.6875%				
Impact Level Probabilities		Impact Scale	Maximum Avoided Cost	Impact scale	Loss (REL)
Low =	60%	Low	$582,000	Low	$ 4,361
Moderate =	30%	Moderate	$1,236,000	Moderate	$ 12,244
High =	10%	High	$2,749,000	High	$ 8,899
				Combined	$25,504
				Approved	□

Carefully review the "Combined" Loss (REL) value. You should be comfortable saying to a Compass User: "Based on your maximum avoided costs, this is what I recommend you prepare to lose from cyber-induced **outages** every year going forward." If the Combined values seem too high or too low, go back to Tab 2 and adjust the baseline probability for outages.

Also review the relative values of the losses from low, moderate, and high events. If they seem disproportionate, go back to Tab 3 and adjust the distribution of impact for **outages.** Once you are satisfied with both the combined values and the distribution, please mark these values "Approved".

4.3 Conditional Probabilities

As mentioned, conditions—including technical controls, policies, architectures, and topologies that users select—can either increase or decrease risks. These increases or decreases are expressed as conditional probabilities that modify the baseline probabilities. This has the effect of making the recommendation for preparedness for a cyberattack more specific to the system under consideration.

To arrive at the conditional probabilities, SMEs provided estimates of adjustment factors to each baseline probability for each possible answer to a condition question. In essence, adjustment factors fine-tune SME estimations for specific characteristics or conditions that apply to the electric system being assessed by the user. For instance, in a condition question about communication protocols, the continued use of older communication protocols that lack security features is likely to increase risk; switching to modern, secure communication protocols decreases risk. The SMEs determined how much these conditions should increase or decrease the baseline and chose an adjustment factor accordingly. These adjustment factors are multiplied with baseline probabilities values to determine the conditional probabilities.

SMEs were allowed to choose adjustment factors between 0.1 and 2.0 of the original baseline probabilities in increments of 0.1. In this way, SMEs could specify that a particular answer would produce a conditional probability from 10% of the baseline probability to 200% of the baseline probability. Table 7 shows the spreadsheet inputs for this.

Table 7. SME Inputs for Adjustment Factors and Conditional Probabilities

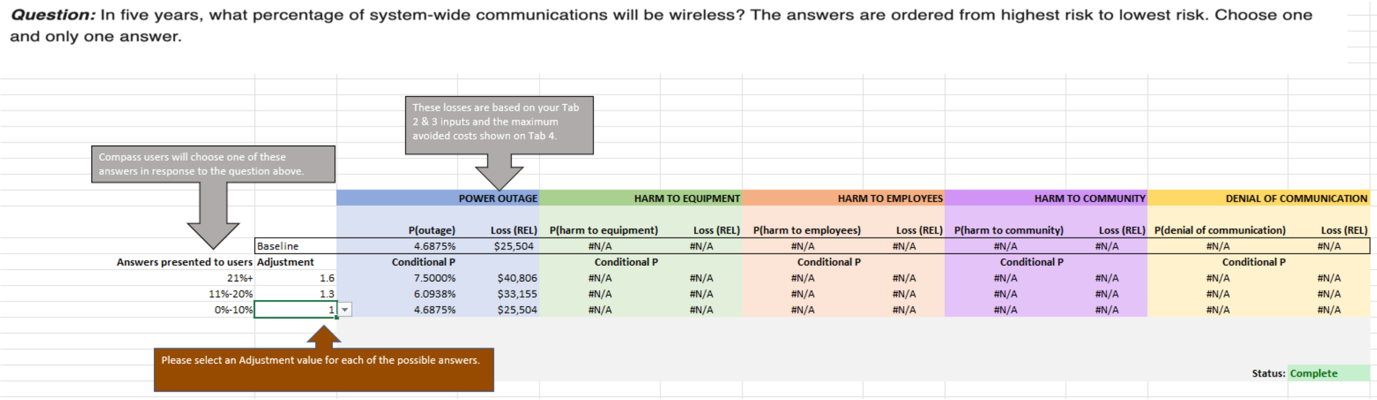

Question: In five years, what percentage of system-wide communications will be wireless? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

Compass users will choose one of these answers in response to the question above.

These losses are based on your Tab 2 & 3 inputs and the maximum avoided costs shown on Tab 4.

		POWER OUTAGE		HARM TO EQUIPMENT		HARM TO EMPLOYEES		HARM TO COMMUNITY		DENIAL OF COMMUNICATION	
		P(outage)	Loss (REL)	P(harm to equipment)	Loss (REL)	P(harm to employees)	Loss (REL)	P(harm to community)	Loss (REL)	P(denial of communication)	Loss (REL)
	Baseline	4.6875%	$25,504	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A
Answers presented to users	Adjustment	Conditional P		Conditional P		Conditional P		Conditional P		Conditional P	
21%+	1.6	7.5000%	$40,806	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A
11%-20%	1.3	6.0938%	$33,155	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A
0%-10%	1	4.6875%	$25,504	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A	#N/A

Please select an Adjustment value for each of the possible answers.

Status: Complete

Note that conditional probabilities are associated with combinations of condition questions posed to users and their answers to those questions. Each possible answer combination has its own adjustment factor and conditional probability. A single user-facing condition question with three possible answers creates three conditions in the back-end analysis (one for each answer). Questions of the form "select one option" contribute a single condition to the analysis. Questions of the form "select all that apply" can contribute conditions numbering from zero to the total number of possible answers supplied. (Appendix E lists all condition questions and their answers.) Figure 2 shows how the front-end data and the back-end data are combined in the Cyber100Compass, as described in Section 3 and Section 4.

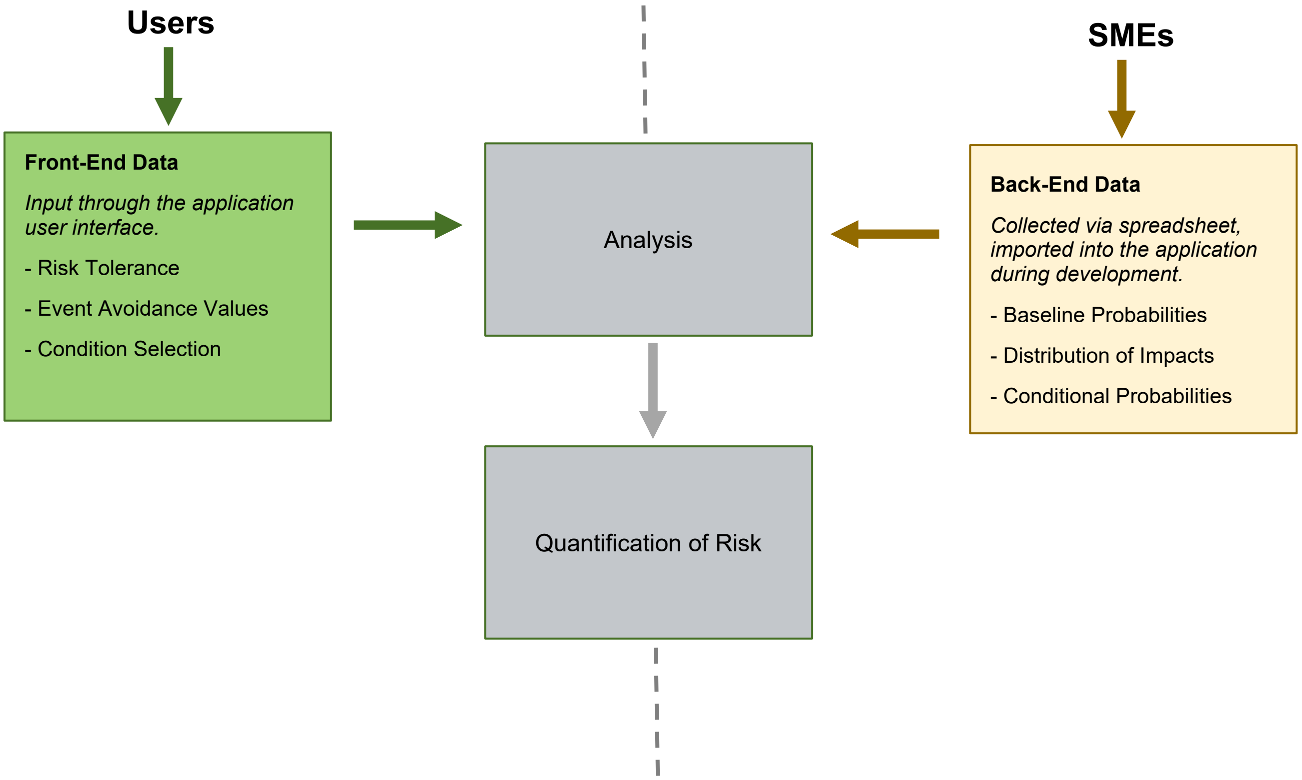

Users
SMEs
Front-End Data
Input through the application user interface.
- Risk Tolerance
- Event Avoidance Values
- Condition Selection
Analysis
Back-End Data
Collected via spreadsheet, imported into the application during development.
- Baseline Probabilities
- Distribution of Impacts
- Conditional Probabilities
Quantification of Risk

Figure 2. Front-end and back-end data and sources

5 Collecting and Aggregating SME Responses

The SMEs' responses to the NREL-facilitated online sessions were collected in SME input spreadsheets. In some cases, the SMEs gathered further input from a committee of knowledgeable colleagues at their organizations. The project eventually obtained four complete sets of SME input.

These four sets of responses were aggregated to produce a single baseline probability for each event, a single set of impact distribution probabilities (the probability of an event occurrence being low-, moderate-, or high-impact) for each event, and a single adjustment factor for each answer to the condition questions.

Two methods were used for the data aggregation: mean and median. A weighted-average method using different weights for each expert was considered but not applied—because no single SME's responses were consistently outliers, weighted aggregation was deemed unnecessary. A comparison between the mean and median responses showed that neither method is noticeably more representative than the other. The mean was selected for being the simplest and best-known method of aggregation. A full comparison between the mean and median is given in Appendix C.

Figure 3 shows a visualization of the baseline probabilities from the four SMEs that contributed data to the proof-of-concept application and the mean and median responses.

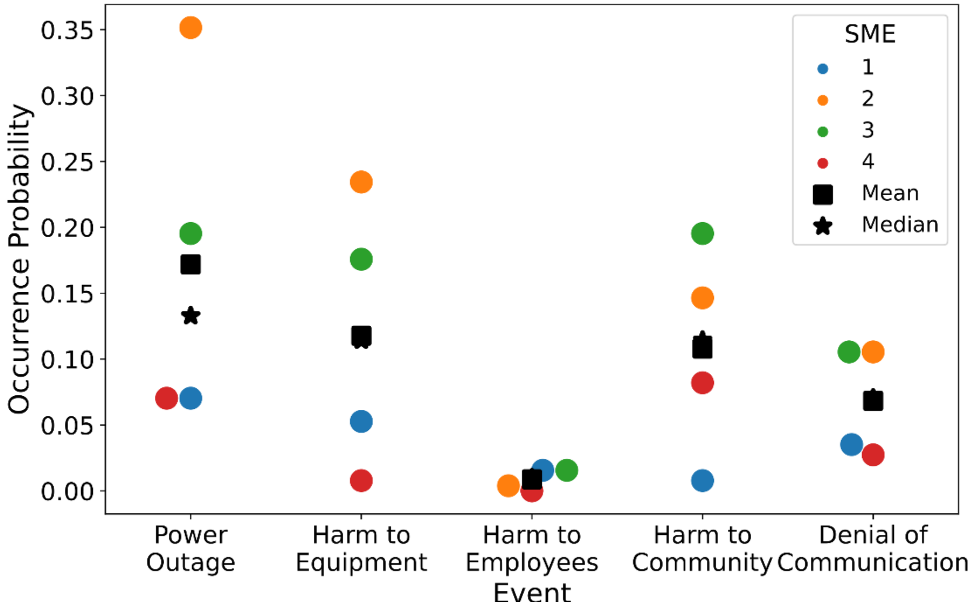

Figure 3. Individual SME responses for baseline probabilities overlaid with aggregate responses in black

The standard deviations serve as a metric for agreement between the SMEs. The standard deviations for the baseline probabilities ranged from 0.008 (for harm to employees) to 0.133 (for power outages). Appendix A provides the standard deviations for all baseline probabilities, distributions of impact, and conditional probabilities.

Note that the SME feedback indicated a flaw with the framing of the "harm to employees" inputs in the SME input spreadsheet. This might have influenced the low values for the harm to employee baseline probabilities shown in Figure 3. The impact of this flaw and how to address similar feedback in follow-on work is discussed further in Section 9. Once the aggregated back-end data are loaded into the application, they are ready for front-end input from the user.

The implementation of the SME data collection presented the project with a number of questions and challenges. The resolutions of these are described as follows:

- Accommodating SME availability:
    - Because SMEs are knowledgeable in both power systems and cybersecurity, they tend to be highly sought (and busy) individuals. The limited time of SMEs suggests that if the proof-of-concept application were to be matured further, SMEs would need to be compensated for their participation.
- Use of distribution of impact:
    - Assessing future cyber risks based on the distribution of impact of cyber events is an innovation of this project. Dividing event losses into levels of low, moderate, and high impact allows more granular specification of loss values (based on user data) and sorting of events into these three levels (based on SME data).
    - An alternative approach that was considered but not pursued would have defined low-, moderate-, and high-impact events as separate events, each with their own baseline probability; however, this would have required much more input from each SME. The SMEs would have been required to choose three baseline probabilities for each event (low, moderate, and high) as well as three conditional probabilities for each answer to each condition question. Not only does this incur significantly more effort, but it requires the SMEs to think in terms of how a particular condition changes the probability for a low-impact event versus a moderate-impact event, etc.
    - Using the distribution of impact values in the Cyber100 Compass reduces the required conditional probability inputs from SMEs by 66%; however, it also means that multiple instances of the same event at different impact levels cannot occur in the same year of the Monte Carlo simulation (see Figure 4 for details). Each event type (e.g., power outage) either occurs or does not occur. The project decided this was an acceptable trade-off that would enable more SME participation.
- Cross-event correlation of adjustment factors:
    - The SME input spreadsheet was created with only one adjustment factor for each condition question answer. In other words, the same adjustment factor applies to each of the five cyber events considered by the Cyber100 Compass.
    - In reality, it is possible (for instance) that a single condition could change the probability of a power outage but not the probability of harm to community, or vice versa. The use of a single adjustment factor for all events was a simplifying assumption to reduce the number of conditional probability inputs required from SMEs by 80%.

    - Note that this 80% reduction in condition-related inputs, together with the 66% reduction obtained by using distributions of impact, resulted in a combined reduction of condition-related inputs of more than 93%. (The remaining condition inputs asked of the SMEs comprised 20% of the original after cross-event correlation, and one-third of that after the application of the distribution of impacts, resulting in approximately 6.67% of the original.) As a result, the overall request made to SMEs for condition inputs was much more reasonable.

- Limits on adjustment factors:
    - Because adjustment factors multiply the baseline probabilities, care must be taken to ensure that the resultant conditional probabilities are not greater than one. This might happen if a SME chooses a baseline probability for a particular event that is close to one as well as conditional adjustment factors on the higher end of the allowable range. (Recall that the highest possible value for an adjustment factor is two.)
    - The spreadsheet used to gather SME data disallows this—if the baseline probability multiplied by the adjustment factor is greater than one, an error message is displayed to the SME; however, this was not an issue for the SMEs who participated in creating back-end data for the proof-of-concept application. A baseline probability greater than 0.5 would mean the SME predicts that there is more than a 50% chance of the utility experiencing that event in a typical year—an extremely pessimistic assessment.

6 Cyber100 Compass Risk Analysis

With back-end data from the SMEs and all front-end data input completed by the user, the user initiates the Cyber100 Compass analysis, which includes the following steps.

6.1 Calculate Conditional Occurrence Probabilities

The Cyber100 Compass assembles a list of applicable conditions based on the user-selected answers to the condition questions.

Note that the number of applicable conditions depends on the selections made by the user. Questions of the format "Choose one and only one answer" include an "Unsure" option, which means that no question-answer combination will be applied. Questions of the format "Choose all that apply" can produce several conditions if multiple answers are selected. Therefore, the number of conditional probabilities that contribute to a single adjusted baseline probability can vary substantially.

Each applicable condition has an associated adjustment factor (A_{C1}, A_{C2}, A_{C3}…) determined by the SMEs. The applicable adjustment factors are multiplied by the baseline probability (P_B) provided by the SMEs to produce a list of conditional probabilities:

$$P_{C1} = A_{C1}P_B \qquad P_{C2} = A_{C2}P_B \qquad P_{C3} = A_{C2}P_B\ldots.$$

The conditional probabilities are then combined with the baseline probabilities using the log odds ratio method to produce an event occurrence probability, p_o. This event occurrence probability combines the SME recommendations for risks with the user plans for their system. Each of the five types of cyber events considered by the application has its own p_o that varies with user input.

6.2 Perform Simulations

The Monte Carlo method is used to produce simulated annual loss values that comprise the output of the application. A Monte Carlo simulation is the process of repeatedly simulating the outcome of an uncertain event or process and then analyzing the outcomes in aggregate to draw inferences about the uncertain event or process (Shonkwiler and Mendivil 2009).

6.3 Determine Ranges of Loss

As explained in Section 3.2, part of the user input required by the Cyber100 Compass is a maximum avoided cost value from each event and at each impact level. In the user interface, these values are referred to as "maximum" values to keep the wording accessible, but they are used as upper bounds of 90% confidence intervals on the loss values. The 90% confidence interval bounds are used to calculate the probability distribution parameters for the losses. For any event, suppose a user provides the upper bound costs UB_L, UB_M, and UB_H for the low-, moderate-, and high-impact losses, respectively. The lower bounds for these losses are generated by the Cyber100 Compass to reduce the amount of user input required. For low-impact losses, the lower bound is set to 1% of the user-provided low-impact upper bound. For moderate-impact losses, the lower bound is set to the low-impact upper bound. For high-impact losses, the lower bound is set to the moderate-impact upper bound. This process is demonstrated in Table 8.

Table 8. Demonstration of How Cyber100 Compass Turns User-Provided Upper Bounds on Loss Values Into Confidence Intervals for Loss Values

Impact Level	User-Provided Upper Bound	Cyber100 Compass-Generated Lower Bound	Cyber100 Compass-Generated 90% Confidence Interval
Low	UB_L	$0.01UB_L$	$[0.01UB_L, UB_L]$
Moderate	UB_M	UB_L	$[UB_L, UB_M]$
High	UB_H	UB_M	$[UB_M, UB_H]$

6.3.1 Monte Carlo Reproducibility, Random Seeds, and Rounds

The results of a Monte Carlo simulation depend on stochasticity, or randomness, and generally involve the use of a random number generator or random draws from probability distributions. Due to this stochasticity, numerical results from a Monte Carlo simulation will be slightly different every time the simulation is executed unless a random seed is used to begin the simulation. A random seed tells the simulation code to use the same process to generate the random numbers every time the simulation is executed, meaning that the Cyber100 Compass will generate the same random numbers and the results will be reproducible.

The Cyber100 Compass sets a random seed by default and provides this value to the users; the seed value can also be set or changed by users. This serves two purposes: The first is reproducibility of the simulation results, and the second is to allow users to iteratively adjust the conditions being applied and view the impact on the simulation results. By setting the random seed before changing the conditions, all changes in the simulation results are due to the updated conditions rather than stochasticity.

There are no concrete rules for how many rounds to execute in a Monte Carlo simulation. Generally, more rounds will provide results that are more stable (that vary less from one simulation to another) and more representative of the underlying uncertain process. The upper limit of rounds that can be run tends to be determined by the computational time and resources available. Within the Cyber100 Compass, the default number of rounds is 100 for performance reasons—users can increase this number until the performance on their systems degrades to unacceptable levels. The developers settled on 100 rounds after performing timed tests of the simulations and assessing the change in the overall simulation results as the number of rounds increased.

6.3.2 Determining Event Occurrence, Impact Level, and Loss Value

The procedure for determining the event occurrence, impact level, and loss value involves several random draws from probability distributions. Event occurrence is modeled with a Bernoulli distribution (Bertsekas and Tsitsiklis 2008):

$$f_o(x_{occur}; p_o) = \begin{cases} p_o & x_{occur} = 1 \\ 1 - p_o & x_{occur} = 0 \end{cases}$$

In this equation, f_o is the probability mass function of the Bernoulli distribution, and p_o is the event occurrence probability that incorporates the effect of any user-applied conditions. This probability mass function means that when the random draw, x_{occur}, from this distribution is equal to 1, the event occurs, and in a single simulated year (one round of the simulation), the event occurs with probability p_o. The outcome of the random draw from a Bernoulli distribution can be thought of as a coin toss with unequal probabilities for the two sides of the coin.

Once an event occurs, a multinoulli, or categorical distribution, models the event impact level (low, moderate, or high) (Murphy 2012):

$$f_v\left(x_{impact} = s \mid (p_L, p_M, p_H)\right) = p_s, \qquad s = (L, M, H)$$

In this equation, f_v is the probability mass function of the multinoulli or categorical distribution; and p_L, p_M, and p_H are the probabilities that the event will be of low, moderate, or high impact, respectively. These probabilities must always sum to 1:

$$p_L + p_M + p_H = 1$$

This probability mass function means that once an event occurs, a random draw, x_{impact}, from this distribution determines the event impact level. The event is of low impact when $x_{impact} = L$, which happens with probability p_L; the event is of moderate impact when $x_{impact} = M$, which happens with probability p_M; and the event is of high impact when $x_{impact} = H$, which happens with probability p_H. The outcome of the random draw from a categorical distribution can be thought of as drawing one of three differently colored items from a bag, with unequal probabilities of drawing each item.

After determining the event impact level, the actual loss value from the event occurrence is simulated as a random draw from the lognormal distribution (NIST 2012a):

$$f_{loss}(x_{loss}) = \frac{1}{x_{loss}\sigma\sqrt{2\pi}} e^{-\frac{(\ln x_{loss} - \mu)^2}{2\sigma^2}}$$

In this equation, f_{loss} is the probability density[1] function of the lognormal distribution, μ is the mean of the distribution, and σ is the standard deviation. The value of μ and σ are calculated by the Cyber100 Compass from the user-provided upper bounds on loss values, which are used to generate 90% confidence intervals on loss values, as shown in Table 8. The equations for μ and σ at each impact level are given in Table 9 (Hubbard and Seiersen 2016, 42).

[1] This is a probability *density* function because the lognormal distribution is continuous. The Bernoulli and multinoulli distributions are discrete—x can only take on certain predefined values—so they are represented with probability *mass* functions.

Table 9. Equations for Calculating the Mean and Standard Deviations of the Lognormal Distributions Used to Model Loss Values

Impact Level	Mean	Standard Deviation
Low	$\mu_L = \dfrac{\ln(0.01UB_L) + \ln UB_L}{2}$	$\sigma_L = \dfrac{\ln UB_L - \ln(0.01UB_L)}{3.28971}$
Moderate	$\mu_M = \dfrac{\ln UB_L + \ln UB_M}{2}$	$\sigma_M = \dfrac{\ln UB_M - \ln UB_L}{3.28971}$
High	$\mu_H = \dfrac{\ln UB_M + \ln UB_H}{2}$	$\sigma_H = \dfrac{\ln UB_H - \ln UB_M}{3.28971}$

6.3.3 Simulation Procedure

Each round of the Monte Carlo simulation represents one typical calendar year. Within a round, each event modeled in the Cyber100 Compass is tested once to determine if it occurs; if the event does occur, then the impact level and the corresponding loss value are determined. If the event does not occur, that event is not tested again until the next round of the simulation. Each event in the Cyber100 Compass can thus occur at most once during one simulation round (representing one calendar year) at a single level of impact. Figure 4 illustrates the overall simulation procedure.

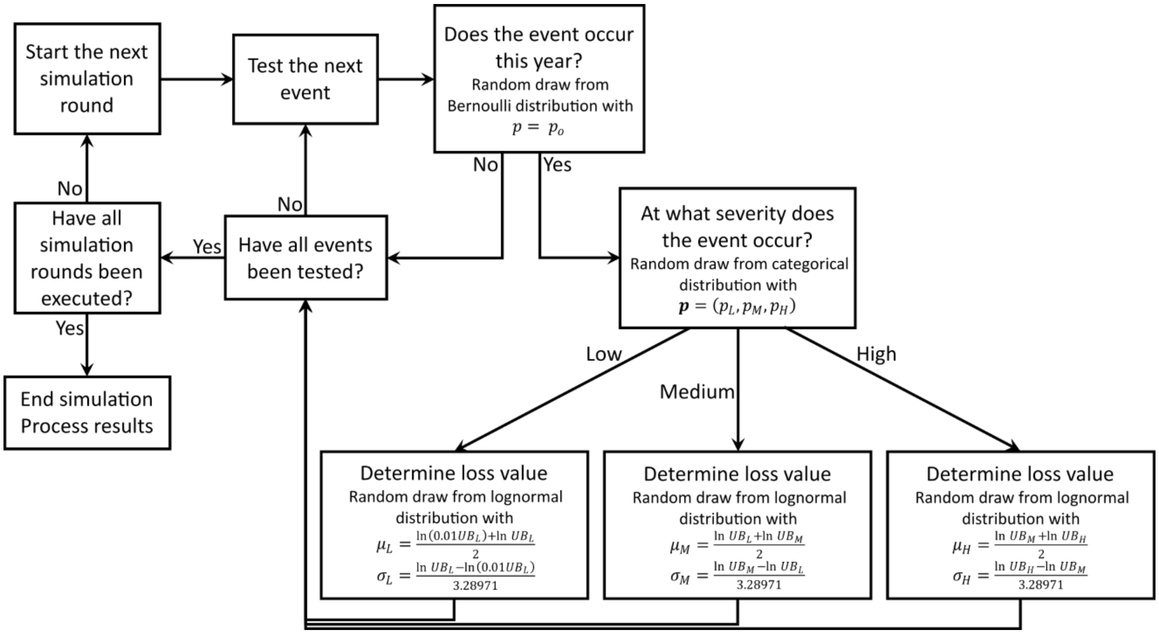
Start the next simulation round
Test the next event
Does the event occur this year?
Random draw from Bernoulli distribution with
$p = p_o$
No
Yes
No
No
Have all simulation rounds been executed?
Yes
Have all events been tested?
At what severity does the event occur?
Random draw from categorical distribution with
$\boldsymbol{p} = (p_L, p_M, p_H)$
Yes
End simulation
Process results
Low
Medium
High
Determine loss value
Random draw from lognormal distribution with
$\mu_L = \frac{\ln(0.01UB_L)+\ln UB_L}{2}$
$\sigma_L = \frac{\ln UB_L - \ln(0.01UB_L)}{3.28971}$
Determine loss value
Random draw from lognormal distribution with
$\mu_M = \frac{\ln UB_L+\ln UB_M}{2}$
$\sigma_M = \frac{\ln UB_M - \ln UB_L}{3.28971}$
Determine loss value
Random draw from lognormal distribution with
$\mu_H = \frac{\ln UB_M+\ln UB_H}{2}$
$\sigma_H = \frac{\ln UB_H - \ln UB_M}{3.28971}$

Figure 4. Flowchart of the Monte Carlo event simulation procedure implemented in the Cyber100 Compass

6.4 Produce Output

The next three figures show example outputs from the proof-of-concept application. These outputs are generated after all required sections have been completed by the users and the Cyber100 Compass has finished the Monte Carlo simulations.

Section 4.1 explained that baseline probabilities are recommendations, rather than predictions, regarding risks faced by utilities. Similarly, the monetized output of the Monte Carlo simulations should be interpreted as recommendations of the type "It is recommended that you prepare your system for annual losses due to cyberattack-induced power outages of $XX per year." As in the SME input spreadsheet, these dollar values are labeled REL in the output graphs.

The REL curve in Figure 5 is derived by sorting the Monte Carlo data into groups according to how many simulations produced losses equal to or greater than the dollar amount on the x-axis. For $1,000, the application compares the number of simulations producing $1,000 or more in losses and divides this by the total number of simulations (100 if the default is used). Likewise for $10,000, $100,000, $1,000,000, and $10,000,000. These percentages are plotted in orange, and a curve is drawn to fit the data points. This is overlaid with the risk tolerance curve created by the user (Section 3.1).

If the REL (orange curve) exceeds the risk tolerance (blue curve), the user might want to adjust elements of their transition plan to reduce risk. If (as shown) the risk tolerance exceeds the REL, the user has confidence that, in the aggregated opinion of the SMEs, their planned transition does not exceed their utility's risk tolerance.

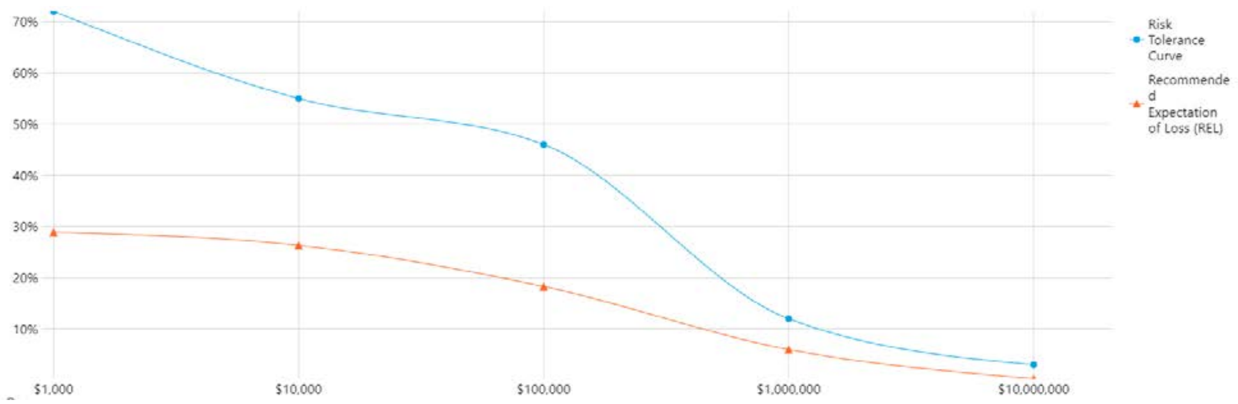

Figure 5. Example output comparing the utility's risk tolerance to REL

Figure 6 shows the average annual monetary losses by event category and level of impact. Figure 7 shows the average annual monetary losses and percentage of overall losses by event type, aggregated over impact levels.

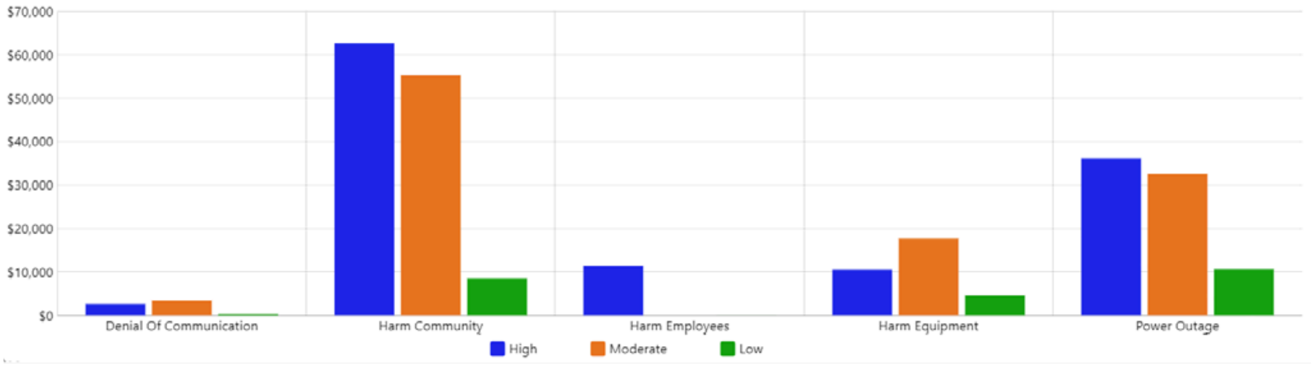

Figure 6. Example output showing average annual monetary impacts by event type and impact level

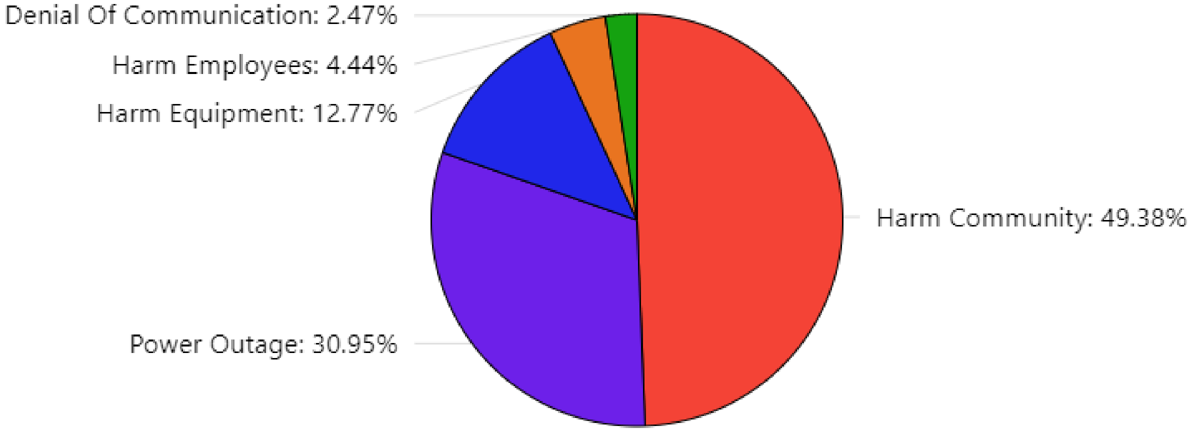

Figure 7. Example output showing average annual monetary losses by event type

7 Known Issues

The Cyber100 Compass project team was able to identify issues that should be addressed in any follow-on work. Additionally, SME feedback on possible future improvements to the Cyber100 Compass are summarized in Appendix B. Also, Gregory Wyss of Sandia National Laboratories produced an independent review of the Cyber100 Compass project (Wyss 2024). This review is included in Appendix F of this report, and some of the observations are referenced throughout this section.

7.1 Limits on SME Participation

The project team understood from the start that SME participation would be a challenge; however, the effort to elicit SME data proved to be even more challenging than anticipated for the following reasons:

- SMEs qualified to provide data need to have familiarity with both electrical systems and the cybersecurity for those systems. These individuals are scarce and (perhaps for that reason) tend to be quite busy. Locating them and arranging data gathering sessions required considerable effort.
- The questions presented to the SMEs regarding conditions must be structured to be "answerable" (i.e., the SME must be able to provide meaningful data on the condition). Decomposing high-level cyber risks into smaller, "answerable" units took considerable effort.
- The questions presented to the SMEs must also be highly structured to ensure that the data gathered from the different SMEs represent the same understanding of the risks and the quantified values placed on them.

For the proof-of-concept application, the SMEs included representatives from utilities, industry service organizations, and consultants. In some cases, SMEs took their input spreadsheets back to their organizations and convened meetings with colleagues to develop input.

Because the project team anticipated limits on SME availability, simplifying assumptions were made to reduce the amount of input required by SMEs. These simplifying assumptions included:

- Applying the same adjustment factor to all cyber events, which reduced SME conditional probability inputs by 80%
- The use of distribution of impact tables, which reduced SME conditional probability inputs by 66%.

The project team developed the SME input spreadsheet to make the SME contributions faster, less confusing, and portable in the sense that SMEs could work on the inputs outside of meetings with NREL and in conjunction with their coworkers.

Still, the project obtained only four complete SME input forms; thus, there were not enough back-end data for the team to feel confident about the results generated by the proof-of-concept application. One reviewer recommended getting at least 30 SMEs to contribute data. The actual number would need to be sufficient to instill confidence not only in the project team but also in the application users.

7.2 Application Implementation Issues

Some issues were observed in the proof-of-concept application after the code was frozen; these would need to be addressed if the application was developed further.

- Unexpected results appear when small baseline probabilities are used. This was observed in the harm to employee event, where the median baseline probability was 0.01. Because this value represents a 1 in 100 probability, and the default number of Monte Carlo runs was 100, this likely indicates that this default needs to be increased to account for small baseline probabilities.
- For the condition "Number of operating entities that will be contributing to generation," the possible answer of 0 should be eliminated.

7.3 Wording for Descriptions of Events

The SMEs pointed out a flaw in the description of the level of impact for harm to employee events. The question asked about the probability of an event that led to physical harm to an employee but specifically did not lead to a hospital visit. The SMEs indicated that if harm to an employee occurred, a hospital visit would be mandatory. This error probably impacted the collected values for that baseline probability.

Errors such as this could be avoided in follow-on work by having a subset of SMEs review descriptions before they are pushed out to the larger groups of SMEs and users. The subset of SMEs could help clarify the text and avoid these kinds of issues.

7.4 Limits of the Mathematical Approach to Conditions

The log odds ratio method used for combining conditions does not capture complex relationships between conditions. As implemented, each condition functions as an independent variable, adjusting the baseline probability without influencing any other condition. In reality, certain combinations of conditions might alter risks in complex ways.

A different mathematical method, the lens method, accounts for such complexities. The lens method requires SMEs to estimate impact costs on combinations of conditions, then it applies logistic regression to those estimates to infer the rules by which the conditions affect each other. The lens method has the advantage of enabling researchers to measure and eliminate inconsistencies in SME input; however, this method requires SMEs to respond to a larger number of inputs (Hubbard and Seiersen 2016, 180–89). Because the project team correctly anticipated that time for SME participation would be at a premium, the log odds ratio was used.

7.5 Limits of Mathematical Approach to Events

The Wyss review of the Cyber100 Compass expresses concern regarding several aspects of how risks were decomposed into the quantification of events. First, Wyss noted that using a Bernoulli distribution (discrete probability distribution of a random variable) excludes the possibility that an event (particularly those with a high likelihood of occurrence) can happen more than once a year. Wyss suggests using a Poisson distribution instead, which would allow the Cyber100 Compass to analyze the frequency of discrete event occurrences over a given time period, rather than a Boolean value expressing whether the event occurred or did not occur.

7.6 “Unsure” As a User Input

Wyss expressed concern about the option for users to select “unsure” as an acceptable answer to the condition questions. This option means that the condition has no impact on the cybersecurity risk, so, in theory, its selection could produce risks either artificially high or artificially low. In practice, however, Wyss argues that the unsure option is more likely to make risks artificially low because it is more likely to be used by utilities with less understanding of cybersecurity.

7.7 Limitations of Calculations of Baseline Probability

Wyss also noted a potential issue in the mathematical approach used to solicit SME input for calculating the baseline probability. As previously discussed, the Cyber100 Compass draws upon the NIST Special Publication 800-30 definitions and assessment scales for determining the likelihood of event occurrence—namely, C= adversary capability, T = adversary targeting, and V = vulnerability severity. Wyss disagrees with the Cyber100 Compass treatment of these factors as independent variables. The multiplication of these factors would be valid only if the probability of each was independent or conditional.

Further, Wyss disagreed with the Cyber100 Compass approach to SME solicitation. First, Wyss noted that the Cyber100 Compass SME elicitation process does not indicate that it follows guidance provided by Hubbard and Seiersen (Hubbard and Seiersen 2016, 133) for ensuring that SME expertise is “calibrated” to reduce quantitative biases and properly considers uncertainties prior to elicitation. In addition, Wyss noted inconsistencies between the SME elicitation best practices and the Cyber100 Compass approach. NIST Special Publication 800-30 uses the definition of the qualitative ordinal values—very high, high, etc.—to help users rate each of the three factors (adversary capability, adversary targeting, vulnerability severity); however, NIST does not ascribe any probabilities to these ratings. The Cyber100 Compass infers probabilities from these ordinal values and asks SMEs to adjust their qualitative answers until the expected loss values align with their expert judgement. Wyss noted that an independent review from an organization such as the National Academies would likely strongly object to this expert elicitation approach as well as the mathematical approach to calculating baseline probability in the current version of the Cyber100 Compass.

7.8 Probability Uncertainty

Wyss noted a limitation regarding the Cyber100 Compass approach to the uncertainty of probabilities. The tool in its current form uses what Wyss called “point estimate values” of computed risks and a ranking of specific scenarios’ risk; however, Wyss stated this method might give users a false sense of confidence when ranking the importance of individual scenarios and making risk mitigation decisions, particularly if the volume of cyberattack scenarios increases in the future and prioritizing mitigations becomes more complex. Wyss noted that a more widely accepted practice and a method to accommodate uncertainty in probabilities is to use confidence intervals to capture uncertainties in results, not discrete point estimates related to the REL.

Further, though there is an unavoidable amount of uncertainty surrounding the most effective cyber risk mitigation measures, Wyss recommended that future versions of the Cyber100 Compass consider and integrate common measures found within resilience mitigation—i.e., look to mitigations that are applicable across a broader range of scenarios, e.g., human error, natural disasters, and malicious cyberattacks, to help compensate for the extremely difficult problem of estimating adversary decision making. Further, incorporating more resilience measures would help with the sparse data problem as well as be more responsive to the wide variability of threats facing evolving electric systems, not only threats from malicious actors.

Additionally, Wyss suggested a complementary risk assessment approach for adversaries with advanced capabilities because these kinds of nation-state or other highly resourced "high-tier" adversaries introduce additional complexities in anticipating adversarial decision making and associated risk management. For this, Wyss suggested looking at established physical security risk management processes, which examine a "high-tier adversary's attack planning process" to understand which mitigations to prioritize based on consequences of successful attacks.

7.9 Academic Controversies About Likelihood of Occurrence

Wyss noted that the most likely and biggest probabilistic uncertainty within the tool is the probability of occurrence (p_o), or the likelihood of a cyberattack that causes physical effects to electric systems. Wyss called this the "problem of quantifying deterrence." Estimating the p_o requires some knowledge or understanding of threat intelligence regarding adversary behavior. That begs the question of how we should consider the likelihood of a cyberattack within a risk assessment such as the Cyber100 Compass. The decision of a malicious actor to initiate an attack and the attackers own cost-benefit risk assessment is poorly understood by defenders and can shift quickly based on world events. Wyss described an academic debate underway between analysts attempting to answer the "problem of quantifying deterrence." To summarize, one side of the debate suggests using Bayesian probability or frequency to quantify deterrence (Hubbard and Seiersen 2016). The other side of the debate suggests that there are probabilistic factors that contribute to the probability of an attack that should be the output of a risk assessment, not an input (Cox 2008). Wyss stated there can be many "hidden" dependencies surrounding the probability of occurrence, making it extremely difficult to accurately quantify $p_{o.}$ Wyss suggested exploring other methods for decomposing these probabilities. One possible approach is to separate p_o into two parts: $p_A \cdot p_{Cn}$, where p_A represents the likelihood that an attack leads to an adverse physical consequence, and p_{Cn} represents the conditional probability of a distinct outcome given a successful attack.

7.10 Back-End Data Anomalies

Where possible, answers for condition questions were structured as progressions, where the first option contributed the most risk and the last option contributed the least risk; however, it was observed that in some cases the adjustment factors provided by the SMEs trended in the other direction—options that the project team thought were most risky were rated least risky and vice versa. It is not clear whether the SMEs disagreed with the project team's assumption about the risk or whether the SMEs were confused by either the question or the adjustment factor input mechanism.

In other cases, SMEs provided answers where the risk was high for the first and last options but lower in between (i.e., the adjustment factors formed a curve rather than a progression.) Again, it was not clear whether the SMEs were confused by the question or had different perspectives on the risks. Errors such as this might be avoided in follow-on work by having SMEs include an explanation for their responses on the SME input form.

8 Possible Follow-On Research

8.1 Validating Application Output

Validation of the application output represents a formidable challenge. Results are meant to be forward-looking, so even if historical data were available, these data would be of limited value. Further progress on the Cyber100 Compass depends on identifying other tools for validation.

Because the application output represents REL based on the user input and the back-end data provided by the SMEs, one approach to validation might measure to what extent the application captures the intent of the SMEs regarding those recommendations.

In this approach, a utility system planner would use the application to perform a Cyber100 Compass assessment. A group of SMEs—either the same group that contributed the back-end data or a group of similarly knowledgeable individuals—would convene to produce their own version of the REL values based on the user input but not using the application. The SMEs would express each element of the output (total loss, loss by cyber event, loss by impact level, etc.) as a range. The application output would then be compared to those ranges, and a score would be produced based on how many of the application-produced values fall within the SME-produced ranges. This process could be repeated with multiple users to understand how well the application aligns with the SMEs' intentions across a range of utility sizes and planned transitions to renewables.

8.2 Evolving the Condition Question Set

The Cyber100 Compass framework was intended to be scalable with regard to conditions. As new conditions are identified as being relevant to risks and answerable by the user, they should be added to the application using the same process as the original conditions: Convene SMEs to provide adjustment factors.

But it is unlikely that the same SMEs would be available to provide these adjustment factors. Would the adjustment factors provided by the new SME group be treated the same as the original adjustment factors, or would they be weighted based on the size of the group (and potentially other factors)? Would the new group also be required to provide baseline probabilities? Would the provenance of the adjustment factors need to be tracked to demonstrate confidence in the data? These questions are embedded in a larger discussion of how an application built on the Cyber100 Compass framework should evolve over time.

If the Cyber100 Compass was released as a user-ready application, it might be advantageous to update it annually or biannually both to expand the condition question set and to refresh the back-end data based on evolving threats, vulnerabilities, and trends. NREL's Regional Energy Deployment System (ReEDS™) already undergoes these sorts of regular updates.

8.3 Incorporating Other Sources of Back-end Data

Although SME data were used in the proof-of-concept application, they were not the only source of data available for the risk assessment. The insurance industry traditionally relies on historical data to assess risks, whereas risks for new technologies and systems are sometimes explored through experiments (Harman, Cortes, and Hill 2018). Future research could explore ways to incorporate historical data and experimental data into the Cyber100 Compass.

8.4 Integration with Other Cybersecurity Tools

The Cyber100 Compass could interoperate with other cybersecurity tools in multiple ways. For example, the output of NREL's Distributed Energy Resource Cybersecurity Framework (DER-CF) could be used to adjust the Cyber100 Compass baseline probabilities. A review of cybersecurity tools originating from within NREL, other national laboratories, academia, and private industry could be performed to identify the best candidates for integration.

9 Conclusion

The Cyber100 Compass represents an attempt to address a scarce data challenge: How to assess risks associated with an electric utility transition to high levels of renewables. This work proceeded by eliciting cyber risk data from SMEs in a format that could be reused in multiple assessments by multiple utilities of various sizes, loads, and generation mixes.

The project successfully developed tools for this type of SME data elicitation, a structured format for this SME data that reduced the amount of condition-related input required by more than 93%, a spreadsheet for collecting the data in that structured format, and processes for combining data from multiple SMEs into a proof-of-concept application.

That proof-of-concept application presents a series of questions to users regarding their organization's appetite for risk, the value they place on avoiding certain types of cyberattack-induced physical events (loss of power, harm to equipment, etc.), and conditions that the system planners expect to be true regarding their systems as they transition to high levels of renewables. Then the user initiates a Monte Carlo simulation, and the application produces quantified risks for the planned transition.

The project successfully developed a set of user-facing questions that collected the three types of front-end data, with supporting text providing background and context for the questions; the analytic engine to combine user data with SME data; and the output graphs and explanatory text. The quantified outputs are REL values—a suggested level of preparedness rather than a prediction of losses.

The SMEs who helped produce the data had many observations about the framework and the application. These are captured and recorded in Appendix B. A number of ideas for follow-on efforts were contributed by both the project team and the SMEs—a prioritized list of these is presented in Appendix A.

The challenge of quantifying risks for systems transitioning to high levels of renewables is embedded in a larger set of challenges involving cybersecurity risk analysis in situations where data are sparse. The Cyber100 Compass framework and the concepts developed in this project might have applications for other sparse data problems facing utilities in the future.

This project represents a first step toward providing system planners with insight about the cyber risks they face as they transition to high levels of renewables. One point of agreement among SMEs, potential users, consultants, and others is that providing this insight is challenging but critical. Whether that is the Cyber100 Compass or another approach, methods to reduce uncertainty about cyber risks in this area will find an eager audience.

Glossary

Term	Definition
Adjustment factor	A value used to adjust the baseline probability estimate from the subject matter expert to correspond to specific electric grid systems being assessed by the user
Baseline probability	The cybersecurity risks present within a user's electric system before any conditions are applied. Also thought of as the level of risk for which a hypothetical "average" utility should prepare
Condition	Any constraint, resource, requirement, control, or other factor that modifies the cybersecurity risks of a system's energy transformation plan. In the Cyber100 Compass, a condition is presented as a series of questions posed to users about their energy transition plan.
Conditional probability	A probability derived by multiplying a baseline probability by an adjustment factor. Also thought of as the level of risk for which a hypothetical "average" utility should prepare given a single condition
Cyber100 Compass	The conceptual basis for the Cyber100 Compass risk analysis
Event	A cyber incident that has a physical impact on operational technology systems and networks—e.g., a power outage, damage to equipment, or the loss of communication leading to the loss of generation
Event occurrence probability	A probability derived from the baseline probability and all applicable conditional probabilities using the log odds ratio. Also thought of as the level of risk for which a hypothetical "average" utility should prepare given all applicable conditions
Log odds ratio	A method of combining the conditional probabilities with a baseline probability to produce the event occurrence probability, p_o
Maximum avoided cost	A monetary value provided by Cyber100 Compass users that describes the highest dollar amount a user estimates the organization could lose from a cyber event in a given year given certain limiting factors. It is provided by users at low-, moderate-, and high-impact levels for each type of Cyber100 Compass cyber event.
Monte Carlo simulation	A mathematical process of repeatedly simulating the outcome of an uncertain event or process and then analyzing the outcomes in aggregate to draw inferences about the uncertain event or process

Term	Definition
Proof-of-concept application	An application created to instantiate the Cyber100 Compass framework, aid in the development and advancement of the framework, and demonstrate the functions of the framework
Random seed	A model parameter within a Monte Carlo simulation that ensures that the same series of pseudo-random numbers are used every time the simulation is executed, ensuring that the results will be reproducible
Recommended expectation of loss	A recommendation from the application regarding how much loss due to a cyberattack the utility should prepare for
Risk tolerance	An organization’s willingness to accept certain levels of risk based on the financial losses that could occur from a cyberattack
Subject atter expert	An individual who is knowledgeable about both cyber risks and utility planning and contributes to the back-end data tables used by the proof-of-concept application. This contribution takes place during the development of the application and must be completed before the application is ready for use.
User	The individual who will use the Cyber100 Compass—e.g., a system planner who needs to assess the cyber risks associated with their plan to transition to high levels of renewable energy

References

Bertsekas, Dimitri, and John Tsitsiklis. 2008. *Introduction to Probability*. Athena Scientific. http://athenasc.com/probbook.html.

CESA. 2023. "Table of 100% Clean Energy States." Clean Energy States Alliance (CESA). 2023. https://www.cesa.org/projects/100-clean-energy-collaborative/guide/table-of-100-clean-energy-states/.

Cox, Louis Anthony (Tony) Jr. 2008. "What's Wrong with Risk Matrices?" *Risk Analysis* 28 (2): 497–512. https://doi.org/10.1111/j.1539-6924.2008.01030.x.

CrowdStrike. 2023. "CrowdStrike 2023 Global Threat Report." CrowdStrike. https://go.crowdstrike.com/rs/281-OBQ-266/images/CrowdStrike2023GlobalThreatReport.pdf.

EIA. 2023. "'How Much of U.S. Energy Consumption and Electricity Generation Comes from Renewable Energy Sources?' Frequently Asked Questions (FAQs)." U.S. Energy Information Administration (EIA). September 28, 2023. https://www.eia.gov/tools/faqs/faq.php?id=92&t=4.

Greenberg, Andy. 2017. "How An Entire Nation Became Russia's Test Lab for Cyberwar." WIRED. June 20, 2017. https://www.wired.com/story/russian-hackers-attack-ukraine/.

———. 2022. "Pipedream Malware: Feds Uncover 'Swiss Army Knife' for Industrial System Hacking | WIRED." *WIRED*, April 13, 2022. https://www.wired.com/story/pipedream-ics-malware/.

Harman, Michael, Luis Cortes, and Raymond Hill. 2018. "Quantifying Test Risk Using Design of Experiments." STAT COE-Report-19-2014. Scientific Test & Analysis Techniques Center of Excellence. https://www.afit.edu/stat/statcoe_files/Quantifying%20Test%20Risk%20Using%20Design%20of%20Experiments%20Rev%201.pdf.

Hubbard, Douglas W., and Richard Seiersen. 2016. *How to Measure Anything in Cybersecurity Risk*. John Wiley & Sons, Inc.

Mathews, Lee. 2017. "NotPetya Ransomware Attack Cost Shipping Giant Maersk Over $200 Million." *Forbes*, August 16, 2017. https://www.forbes.com/sites/leemathews/2017/08/16/notpetya-ransomware-attack-cost-shipping-giant-maersk-over-200-million/?sh=35dc62284f9a.

Murphy, Kevin. 2012. *Machine Learning: A Probabilistic Perspective*. MIT Press. https://probml.github.io/pml-book/book0.html.

NIST. 2012a. "1.3.6.6.9. Lognormal Distribution." In *NIST/SEMATECH e-Handbook of Statistical Methods*. National Institute of Standards and Technology (NIST). https://doi.org/10.18434/M32189.

———. 2012b. "Guide for Conducting Risk Assessments." NIST Special Publication (SP) 800-30 Rev. 1. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-30r1.

NIST, National Institute of Standards and Technology. n.d. "Glossary." Computer Security Resource Center. Accessed October 15, 2023. https://csrc.nist.gov/glossary/term/predisposing_condition.

NREL. 2022. "Valuing Resilience in Electricity Systems." NREL/FS-7A40-74637. National Renewable Energy Laboratory (NREL). https://www.nrel.gov/docs/fy19osti/74673.pdf.

Sheehan, Barry, Finbarr Murphy, Arash N. Kia, and Ronan Kiely. 2021. "A Quantitative Bow-Tie Cyber Risk Classification and Assessment Framework." *Journal of Risk Research* 24 (12): 1619–38. https://doi.org/10.1080/13669877.2021.1900337.

Shonkwiler, Ronald W., and Franklin Mendivil. 2009. *Explorations in Monte Carlo Methods*. Undergraduate Texts in Mathematics. New York, NY: Springer. https://doi.org/10.1007/978-0-387-87837-9.

Wood, Kimberly. 2023. "Cybersecurity Policy Responses to the Colonial Pipeline Ransomware Attack." *The Georgetown Environmental Law Review*, March 7, 2023. https://www.law.georgetown.edu/environmental-law-review/blog/cybersecurity-policy-responses-to-the-colonial-pipeline-ransomware-attack/.

Wyss, Gregory. 2024. "Independent Review of the Proof-of-Concept Cyber100 Compass Cybersecurity Risk Tool." SAND2024-03315R. Sandia National Laboratories.

Zetter, Kim. 2014. *Countdown to Zero Day*. New York: Crown Publishers.

Bibliography

Agrafiotis, Ioannis, Jason R. C. Nurse, Michael Goldsmith, Sadie Creese, and David Upton. 2018. "A Taxonomy of Cyber-Harms: Defining the Impacts of Cyber-Attacks and Understanding How They Propagate." *Journal of Cybersecurity* 4 (tyy006). https://doi.org/10.1093/cybsec/tyy006.

Bartol, Nadya. 2015. *Cyber Supply Chain Risk Management for Utilities—Roadmap for Implementation*. Washington, D.C.: Utilities Telecom Council. https://utc.org/wp-content/uploads/2018/02/SupplyChain2015-2.pdf.

Bodeau, Deborah J., and Catherine D. McCollum. 2018. *System-of-Systems Threat Modeling*. McClean, VA: MITRE, Homeland Security Systems Engineering & Development Institute. Case Number 18-1631/DHS reference number 16-J-00184-07. https://www.mitre.org/sites/default/files/publications/pr_18-1631-ngci-system-of-systems-threat-model.pdf.

Carter, Cedric, Christine Lai, Nicholas Jacobs, Shamina Hossain-McKenzie, Patricia Cordeiro,Ifeoma Onunkwo, and Jay Tillay Johnson. 2017. *Cyber Security Primer for DER Vendors Aggregators and Grid Operators*. Albuquerque, NM: Sandia National Laboratories. SAND-2017-13113. https://doi.org/10.2172/1761987.

Chiprianov, Vanea, Laurent Gallon, Manuel Munier, Philippe Aniorte, and Vincent Lalanne. 2014. "The Systems-of-Systems Challenge in Security Engineering." In *Journées Nationales 2014 Du GDR Génie de La Programmation et Du Logiciel (GPL'2014)*, 163–66. Paris, France. http://munier.perso.univ-pau.fr/en/publication/2014/2014-gpl/.

Dahmann, Judish. n.d. "System of Systems Characterization and Types." North Atlantic Treaty Organization Science and Technology Organization. STO-EN-SCI-276. https://www.sto.nato.int/publications/STO%20Educational%20Notes/STO-EN-SCI-276/EN-SCI-276-01.pdf.

U.S. Department of Defense (DOD). 2010. "Systems Engineering Guide for Systems of Systems: Summary." Washington, DC: Director of Systems Engineering, Office of the Director, Defense Research and Engineering. https://www.acqnotes.com/Attachments/Summary%20-%20Systems%20Engineering%20Guide%20for%20Systems%20of%20Systems.pdf.

Goff, Ed, Cliff Glantz, and Rebecca Massello. 2014. "Cybersecurity Procurement Language for Energy Delivery Systems." In *Proceedings of the 9th Annual Cyber and Information Security Research Conference (CISR '14)*, 77–79. New York, NY: Association for Computing Machinery. https://doi.org/10.1145/2602087.2602097.

Institute of Electrical and Electronics Engineers (IEEE). 2012. IEEE Std 1366-2012 (Revision of IEEE Std 1366-2003)—IEEE Guide for Electric Power Distribution Reliability Indices. New York, NY. https://doi.org/10.1109/IEEESTD.2012.6209381.

Interruption Cost Estimate (ICE) Calculator. 2023. U.S. Department of Energy Office of Electricity, Lawrence Berkeley National Laboratory, and Nexant. Accessed June 8, 2023. https://icecalculator.com/home.

Johnson, Jay. 2017. *Roadmap for Photovoltaic Cyber Security*. Albuquerque, NM: Sandia National Laboratories. SAND2017-13262, 1782667, 668568. https://doi.org/10.2172/1782667.

Ki-Aries, Duncan, Shamal Faily, Huseyin Dogan, and Christopher Williams. 2018. "System of Systems Characterisation Assisting Security Risk Assessment." In *Proceedings of the 2018 13th Annual Conference on System of Systems Engineering (SoSE)*, 485–92. https://doi.org/10.1109/SYSOSE.2018.8428765.

Kniesner, Thomas J., and W. Kip Viscusi. 2019. "The Value of a Statistical Life." *Oxford Research Encyclopedia of Economics and Finance*. Vanderbilt Law Research Paper No. 19-15. https://doi.org/10.2139/ssrn.3379967.

MITRE. 2017. "Treating Systems of Systems as Systems." November 28, 2017. https://www.mitre.org/publications/systems-engineering-guide/enterprise-engineering/systems-of-systems/treating-systems-of.

Muhlberg, Byron. 2020. "U.S. Critical Infrastructure Victim of Ransomware Attack." *CPO Magazine*. March 5, 2020. https://www.cpomagazine.com/cyber-security/u-s-critical-infrastructure-victim-of-ransomware-attack/.

Narang, David, Peter Schwartz, Steve Widergren, Sigifredo Gonzalez, S. Alam, Theodore Bohn, Yaosuo Xue et al. 2021. *GMLC Survey of Distributed Energy Resource Interconnection and Interoperability Standards*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5D00-77497, 1823018, MainId: 27433. https://doi.org/10.2172/1823018.

National Conference of State Legislators. 2021. "State Renewable Portfolio Standards and Goals." August 13, 2021. https://www.ncsl.org/research/energy/renewable-portfolio-standards.aspx.

National Institute of Standards and Technology (NIST). 2020. *NIST Special Publication 800-53 Revision 5: Security and Privacy Controls for Information Systems and Organizations*. Gaithersburg, MD. https://doi.org/10.6028/NIST.SP.800-53r5.

National Renewable Energy Laboratory (NREL). 2023. "Customer Damage Function Calculator." Accessed June 8, 2023. https://cdfc.nrel.gov.

Occupational Safety and Health Administration (OSHA). 2023. "Estimated Costs of Occupational Injuries and Illnesses and Estimated Impact on a Company's Profitability Worksheet." U.S. Department of Labor. Accessed June 8, 2023. https://www.osha.gov/safetypays/estimator.

Onunkwo, Ifeoma. 2020. *Recommendations for Data-in-Transit Requirements for Securing DER Communications*. Albuquerque, NM: Sandia National Laboratories. SAND2020-12704. https://doi.org/10.2172/1813646.

Pinar, Ali, Thomas Tarman, Laura Painton Swiler, Jared Gearhart, Derek Hart, Eric Vugrin, Gerardo Cruz, et al. 2021. *Science and Engineering of Cybersecurity by Uncertainty Quantification and Rigorous Experimentation (SECURE) (Final Report)*. Albuquerque, NM: Sandia National Laboratories. SAND-2021-11719. https://doi.org/10.2172/1821322.

Reaves, Bradley, and Thomas Morris. 2012. "Analysis and Mitigation of Vulnerabilities in Short-Range Wireless Communications for Industrial Control Systems." *International Journal of Critical Infrastructure Protection 5* (3): 154–74. https://doi.org/10.1016/j.ijcip.2012.10.001.

Shea, Daniel. 2020. *Cybersecurity and the Electric Grid: The State Role in Protecting Critical Infrastructure*. Washington, D.C.: National Conference of State Legislatures. Accessed June 8, 2023. https://www.ncsl.org/energy/cybersecurity-and-the-electric-grid-the-state-role-in-protecting-critical-infrastructure.

Stouffer, Keith, Victoria Pillitteri, Suzanne Lightman, Marshall Abrams , and Adam Hahn. 2015. *NIST Special Publication 800-82, Revision 2: Guide to Industrial Control Systems (ICS) Security*. Gaithersburg, MD: National Institute of Standards and Technology. http://dx.doi.org/10.6028/NIST.SP.800-82r2.

Sundararajan, Aditya, Aniket Chavan, Danish Saleem, and Arif I. Sarwat. 2018. "A Survey of Protocol-Level Challenges and Solutions for Distributed Energy Resource Cyber-Physical Security." *Energies* 11 (9): 2360. https://doi.org/10.3390/en11092360.

SunSpec Alliance. 2018. "Cybersecurity Webinar: Securing California Rule 21 Networks." December 13, 2018. https://sunspec.org/cybersecurity-webinar-securing-california-rule-21-networks/.

University of California, Los Angeles (UCLA). 2019. "Progress Toward 100% Clean Energy in Cities and States Across the U.S." UCLA Luskin Center for Innovation. https://innovation.luskin.ucla.edu/wp-content/uploads/2019/11/100-Clean-Energy-Progress-Report-UCLA-2.pdf.

U.S. Department of Energy (DOE). 2022. *Cybersecurity Considerations for Distributed Energy Resources on the U.S. Electric Grid*. Washington, D.C.: U.S. Department of Energy Office of Cybersecurity, Energy Security, and Emergency Response and Office of Energy Efficiency and Renewable Energy. https://www.energy.gov/sites/default/files/2022-10/Cybersecurity%20Considerations%20for%20Distributed%20Energy%20Resources%20on%20the%20U.S.%20Electric%20Grid.pdf.

U.S. Department of Homeland Security (DHS). 2009. *Cyber Security Procurement Language for Control Systems*. Washington, D.C.: DHS Control Systems Security Program, National Cyber Security Division. https://www.cisa.gov/sites/default/files/2023-01/Procurement_Language_Rev4_100809_S508C.pdf.

Wang, Weikang, Kaiqi Sun, Chujie Zeng, Chang Chen, Wei Qiu, Shutang You, and Yilu Liu. 2021. "Information and Communication Infrastructures in Modern Wide-Area Systems." In *Wide Area Power Systems Stability, Protection, and Security*, edited by Hassan Haes Alhelou, Almoataz Y. Abdelaziz, and Pierluigi Siano, 71–104. Power Systems. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-54275-7_3.

Appendix A: Next Phases

From its current state of development, work on the Cyber100 Compass could proceed in several directions. Following is a recommended phased approach to further development. The phases could be done in a different sequence, but this sequence is recommended by the project team.

Phase 1: Advance the proof-of-concept application to a user-ready application. This would require recruiting a test set of system planners to complete the Cyber100 Compass analysis and provide feedback on the clarity of the questions, the difficulty in gathering the required data, the usability of the application interface, and the value of the results. The feedback would be used to modify the application and create new questions that could be added to it. These questions could focus on the relative risks of different renewable technologies (wind, solar, etc.), the application of cyber-informed engineering or consequence-driven cyber-informed engineering, or the evolution of the threat landscape. The subject matter expert (SME) input spreadsheet would be updated based on this feedback.

Once completed, another group of SMEs would be convened to provide more data using the input spreadsheet. This would be loaded as back-end data to the application. Then the SMEs would reconvene to review a number of example user data sets to validate that the application outputs fall within an acceptable range corresponding to the recommended expectation-of-loss values.

Phase 2: Create a web-based version of the application. The current proof-of-concept application is a stand-alone application written in Python and run on the system planner's desktop; however, it was written using libraries that will support porting to a web-based application, which offers several advantages, including ease of updating (e.g., adding new condition questions or correcting errors), the inclusion of back-end data from proprietary sources, easier alignment with the SME input form (by integrating the functionality of the current SME input spreadsheet into the web application), and adjusting the baseline probabilities based on input from NREL's Distributed Energy Resource Cybersecurity Framework (DER-CF).

Phase 3: Explore improvements to the framework algorithm. The limits of the log odds ratio method are discussed in Section 7.2. Other methods might be useful for capturing complex relationships between conditions (e.g., the lens method) or measuring confidence in the output. Changing the methods used on the back end would be a significant pivot in the development of the Cyber100 Compass, but it might be warranted if there is substantial value in other methods.

Phase 4: Explore the incorporation of risk data from other sources. SMEs are not the only possible source of cyber risk data. Historical risk data or data generated through simulations might help improve the application. Both pose challenges, including the integration itself (e.g., how are historical or experimental data weighted relative to SME data?) and validation; however, given the potential of these data sources, eventually, both should be explored.

Appendix B: Subject Matter Expert Engagement

The National Renewable Energy Laboratory (NREL) engaged with subject matter experts (SMEs)—those who have expertise in the fields of power systems, cybersecurity, and risk management—from a variety of organizations to contribute to and create data tables for computing the back-end probabilities to make the Cyber100 Compass run. The SMEs included representatives from:

- American Council on Renewable Energy
- American Public Power Association
- Black & Veatch
- CIP Corps
- National Rural Electric Cooperative Association
- Platte River Power Authority
- Southern California Edison.

These individuals were invited to a series of facilitated discussions and working sessions where they were briefed on the Cyber100 Compass and guided through an exercise to fill out information for the SME input tool. Although not all SMEs contributed to the final data set, each engagement provided valuable feedback and insights.

During these discussions, the SMEs were asked to complete an NREL-created spreadsheet to capture three data inputs to create the back-end probabilities: baseline probabilities, distribution of impact, and conditional probabilities.

The SMEs were also asked to provide feedback on their impressions of the risk assessment approach, the definitions and categorizations of events and conditions, and their opinions about the Cyber100 Compass user experience developed by the Cyber100 Compass project team.

This section summarizes the feedback received from the SMEs as they completed the input spreadsheet and learned about the tool's purpose and approach to evaluating future cybersecurity risks to clean energy transitions.

This section organizes the feedback and responses from the SMEs into two main categories: (1) feedback related to the SME input process and user experience and (2) considerations for improving Cyber100 Compass concepts. Several key themes emerged in both categories.

Several experts agreed that the description of the tool and the engagement sessions were valuable and that the information was presented in a way that made sense. Some SMEs said that their own organizations would benefit from the tool, and others stated that the tool would have been helpful when they first started planning their clean energy transitions several years ago.

Feedback Related to SME Input and User Experience

Participants discussed a range of challenges, considerations, and potential improvements to the process of gaining the needed SME inputs for the back end and the user inputs for the front end of the application. In several instances, this feedback was incorporated into the design or content improvements to the tool. In other cases, the feedback was recorded for consideration in future iterations of the application.

SME Engagement and Other Stakeholders

Several SMEs expressed concern that they did not feel qualified to answer with certainty or felt that the requested inputs were outside their roles or areas of expertise. Others noted that their organizations had not experienced cyber incidents like those described in the tool and suggested that NREL reach out to organizations that had experienced related breaches. The SMEs also suggested engaging with financial experts and/or risk analysts as valuable additional invitees to improve the inputs and feedback received, given that the Cyber100 Compass relies on monetary values to estimate the impacts of different cyber events. Participants also noted that it would be helpful for the SMEs to engage with system engineers because they would have a better understanding of the costs for other things that would be relevant to cyber events that cause damage to electrical equipment, such as building a new substation, building new electric lines, or procuring new transformers. Some SMEs also suggested consulting with the North American Electric Reliability Corporation (NERC), whereas others thought NERC's focus on transmission and bulk power systems would be out of scope for this tool. Last, the SMEs recommended engaging with SMEs in artificial intelligence to account for trends and current events in that field as well.

Participants also noted some confusion between the Cyber100 Compass application itself and the SME input tool, and how one related to the other. Others noted that the SME and user experience would be improved by trimming or breaking up blocks of text to make it easier to follow the instructions. The SMEs cautioned that the availability of staff and resources needed to collect the user inputs could substantially vary based on the size of the utility.

The SMEs also conveyed the need for the Cyber100 Compass to be widely broadcast so that potential users know about the tool and can find it. The SMEs proposed several industry conferences to help NREL's project team market the solution, including the American Public Power Association, the National Rural Electric Cooperative Association, NSI Industries, and the Multi-State Information Sharing and Analysis Center.

Who Should or Could Use Cyber100 Compass

Participants suggested several additional stakeholders beyond the original intended audience that might benefit from the Cyber100 Compass risk assessments. The SMEs mentioned the insurance industry several times—both as an additional source of data for the Cyber100 Compass and as interested parties and users of the tool to validate the insurance risk estimates for operational technology (OT) systems. The SMEs also noted the value of the Cyber100 Compass as a tool to relatively quickly and easily engage with senior executive personnel in conversations around cyber risks. According to the SMEs, it is likely that financial and system design teams would also find the tool useful. For a small team of security personnel at an organization, it is possible that the tool could be used as a good starting point for developing a risk assessment. In general, consensus seemed to be that the results of the Cyber100 Compass risk assessments would be most useful if individuals across an organization were involved in the user input process.

Considerations for Improving Cyber100 Compass Concepts

Participants offered valuable insights into how NREL researchers could improve the concepts, approaches, and assumptions found in the application. This section summarizes the feedback related to the core concepts of the Cyber100 Compass.

Measurement Factors Not Accounted for in the Cyber100 Compass

The SMEs made several comments related to the level of abstraction taken by the Cyber100 Compass and how potentially important details and nuances could be lost as a result. The SMEs shared questions and concerns that the Cyber100 Compass expresses the impacts of cyber events in dollars. Several SMEs expressed concern over using monetary values for all impacts and whether that adequately captures safety and the utility's role in protecting people and the public. SMEs urged the project team to consider how different audiences might react differently to the financialization of impacts, including safety concerns. Experts also expressed concerns over how a cyber attack could bring down the grid and have downstream affects like disrupting a hospital or start a forest fire. Some SMEs recommended translating dollar amounts to qualitative assessments for issues such as safety and downstream effects.

Many SMEs noted that it was not clear whether the tool accounted for the size of generation of different utilities and that perhaps this could be considered and could be more specific in future versions of the tool—for example, the risks for an investor-owned utility and the desire for adversaries to target their systems could substantially vary between a rural cooperative or a municipal utility. The SMEs also raised questions about the conditions that asked users about the criticality of loads. The SMEs noted that some critical loads might be served only by certain substations, and that if users answered more generally, this could exclude the specific nature of more sensitive loads. The SMEs also recommended accounting for the variations in the cybersecurity maturity of different vendors that might interconnect with a user's system.

The SMEs noted that conditions related to the segmentation of communications could be expanded to include considerations related to segmenting the substation architecture. The SMEs referenced a joint report by NERC and the Institute of Electrical and Electronics Engineers on segmentation that recommended segmenting substations so that an attack on one substation would not give attackers unfettered access to all substations. Additionally, the SMEs suggested including condition questions related to redundancy to increase cyber resilience.

Current and Future Threat Landscape

Many SMEs asked questions and gave feedback on current events and related cybersecurity risks stemming from geopolitical tensions.

Several SMEs discussed the challenges in predicting changes in the operating environment and future risks from tensions in international relations. For instance, supply chain risks were mentioned as being based on nation state dependencies. Utilities need to make purchase choices, and many components, such as source code like those for electric vehicles, come from China. What happens to cyber risks if international relations with China worsen? The SMEs expressed the desire for the Cyber100 Compass to account for the interdependencies of the supply chain and geopolitical risks. Adversaries are known to plan decades in advance, so the SMEs believed there should be a way to capture this issue. Other SMEs stated that this might be outside the purview of the tool. Some SMEs expressed concern that Cyber100 Compass was trying to complete a difficult task by capturing systems of interests, vulnerabilities, and relevant threat actors, which are imperfect and challenging proxies for threat intelligence.

The SMEs were intrigued by the absence of two kinds of events from the Cyber100 Compass: ransomware and physical attacks. The SMEs expressed an opinion that it would be more likely for threat actors to cause an outage or deny communications via a physical attack because that is easier than a cyberattack.

The SMEs also raised questions about whether the events described in the Cyber100 Compass could be mutually exclusive or whether the tool would consider multiple event impacts simultaneously occurring. Further, the SMEs added that a low impact level is unlikely for threat actors that want to cause harm. The SMEs noted a lack of desire to accidentally disrupt OT systems. According to the SMEs' experiences, these systems are more likely to be old, outdated, and more vulnerable; however, with the help of technologies such as distributed energy resource management systems, operators will have the ability to dispatch resources, so it will be important to account for technology changes in the future with regard to wind, solar, and other distributed energy resources, as opposed to the substations and infrastructure we have today.

Additionally, SMEs from smaller enterprises were skeptical that threat actors such as a nation state would target small organizations without prioritizing attacks on large utilities with millions, not thousands, of customers.

New Conditions Suggested by DOE SMEs

During the project, the NREL team engaged with SMEs from the U.S. Department of Energy. One point expressed in these conversations was to include condition questions that focused on individual renewable technologies (wind, solar, etc.). The current condition questions in the proof-of-concept application treat all renewable technologies the same with regard to risk. Breaking out each renewable technology would require generating relative risk impacts for each technology.

Other possible condition questions discussed include those based on the application of cyber-informed engineering and consequence-driven cyber-informed engineering.

The comments, suggestions, and perspectives captured in these SME engagements provided valuable feedback that was incorporated to improve the Cyber100 Compass and offered new ideas for future updates and iterations of the application. NREL's project team is grateful for each expert's perspectives and participation in these engagement sessions and their data inputs to the Cyber100 Compass.

Appendix C: Subject Matter Expert Data Analysis

The project team tested and compared two methods of calculating aggregate SME responses: mean and median. No single SME's responses were consistently outliers, so an expert-weighted aggregation method was not applied.

There is very little difference between the aggregation methods, and it is unlikely that the methods will produce substantially different simulation results. We recommend the mean as the aggregation method—it is the simplest, best-known method, and it is likely to be the easiest to justify.

The project team also recommends the mean for the conditional adjustment factors. There was little to no difference between the mean and median adjustment factor responses from the SMEs.

Occurrence and Distribution of Impact Probabilities

Table C-1 shows the means, medians, and standard deviations of the baseline probabilities for each cyber event. Table C-2 shows the means, medians, and standard deviations for the distribution of impact for each cyber event.

Table C-1. Baseline Probabilities and Their Means, Medians, and Standard Deviations

Event	Mean	Median	Standard Deviation
Denial of communication	0.068	0.070	0.043
Harm to community	0.108	0.114	0.081
Harm to employees	0.009	0.010	0.008
Harm to equipment	0.118	0.114	0.105
Power outage	0.172	0.133	0.133

Table C-2. Distributions of Impact and Their Means, Medians, and Standard Deviations

Event	Impact Level	Mean	Median	Standard Deviation
Denial of communication	Low	0.500	0.600	0.283
	Moderate	0.388	0.300	0.214
	High	0.112	0.100	0.075
Harm to community	Low	0.487	0.550	0.239
	Moderate	0.413	0.375	0.210
	High	0.100	0.100	0.041
Harm to employees	Low	0.593	0.750	0.393
	Moderate	0.362	0.200	0.403
	High	0.045	0.050	0.010
Harm to equipment	Low	0.575	0.625	0.240
	Moderate	0.342	0.300	0.187
	High	0.083	0.075	0.054
Power outage	Low	0.562	0.575	0.320
	Moderate	0.320	0.350	0.214
	High	0.118	0.075	0.126

Figure C-1 shows the individual SME responses for the baseline probabilities overlaid with aggregated probabilities in black. Overall, neither aggregated method (mean and median) is noticeably more representative than the other. Figure C-2 shows the individual SME response for the distributions of impact. Figure C-3 compares the means and medians for the distributions of impact. Both methods result in similar trends across impact levels. Neither aggregation method appears to be more representative than the other.

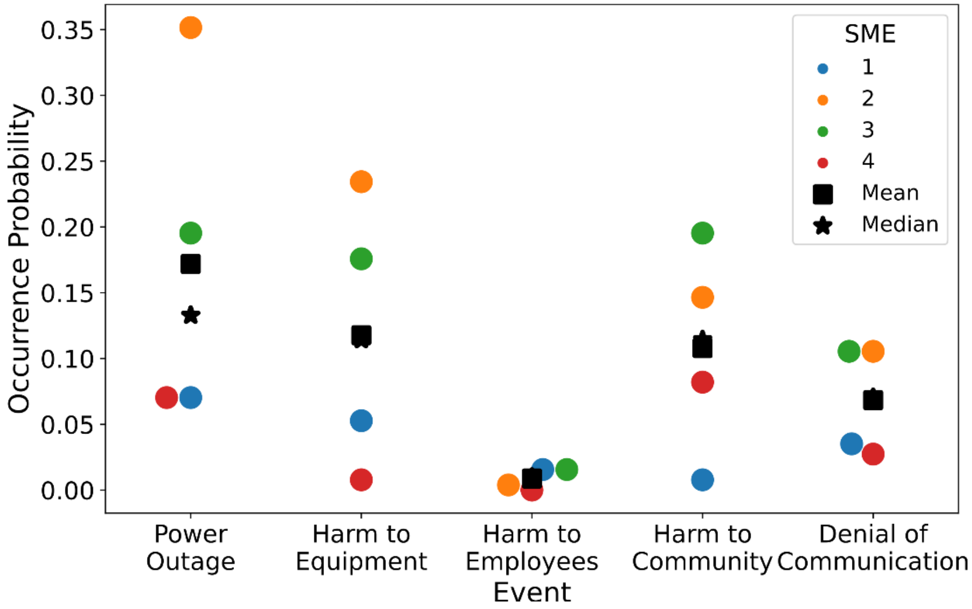

Figure C-1. Individual SME responses for the baseline probabilities

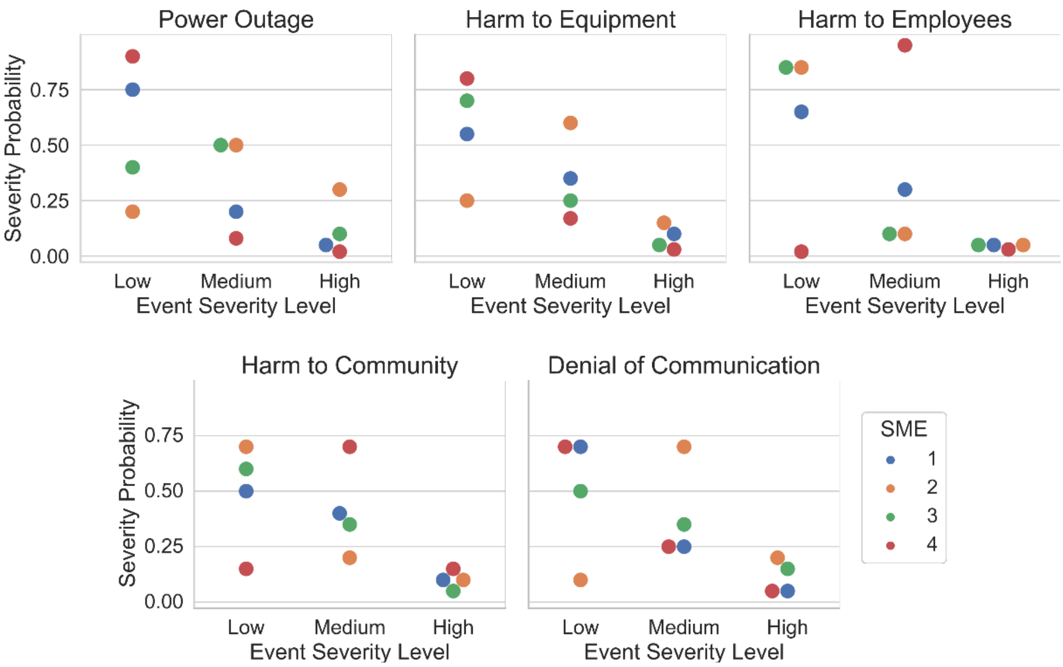

Figure C-2. Individual SME responses for the distributions of impact

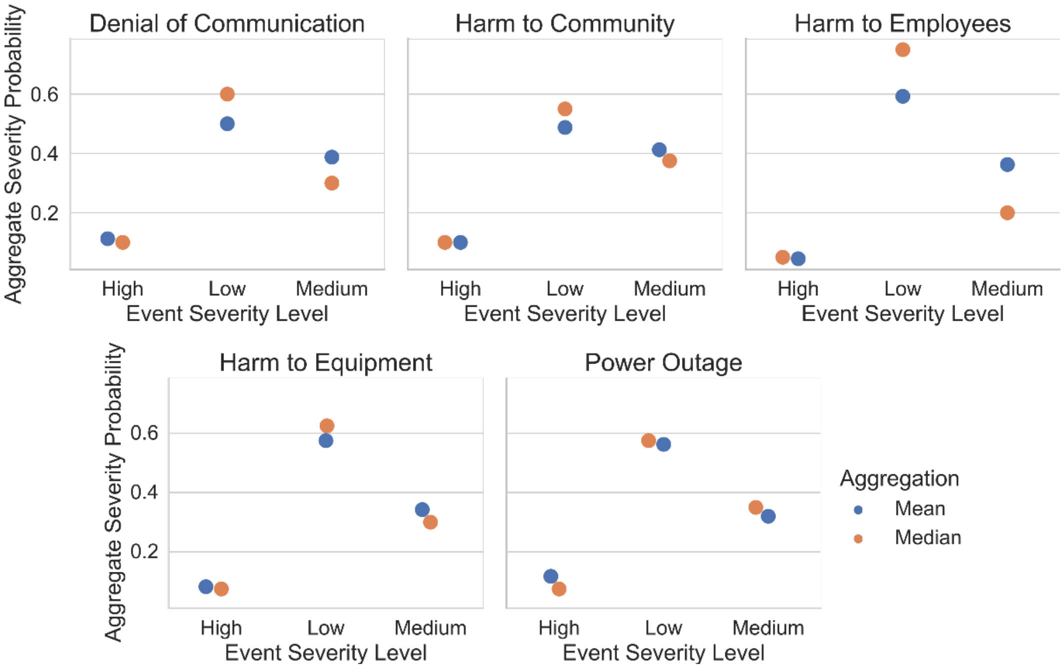

Figure C-3. Comparisons of the methods for aggregating the distributions of impact

Conditional Probability Adjustment Factors

As discussed in Section 6.1, the conditional probabilities are the product of the baseline probabilities and the adjustment factors. Table C-3 shows the means, medians, and standard deviations for all adjustment factors based on the SME inputs. (Note: Some of the questions and answers in this table have been edited. See Appendix E for the complete text.)

Table C-3. Adjustment Factors and Their Means, Medians, and Standard Deviations

Question	Answer	Mean	Median	Standard Deviation
Percentage of generation not operated by utility At the end of five years, what percentage of power on your grid do you expect to be generated by sources not operated by the utility?	76%–100%	1.850	1.90	0.170
	51%–75%	1.500	1.50	0.363
	26%–50%	1.275	1.15	0.455
	0%–25%	0.600	0.65	0.332
Largest nonutility operating entity At the end of five years, what percentage of power capacity (or in the case of storage, energy capacity) on your grid will be supplied by the largest non-utility operating entity?	76%–100%	1.875	2.00	0.222
	51%–75%	1.600	1.60	0.299
	26%–50%	1.375	1.25	0.400
	0%–25%	0.700	0.85	0.377
Percentage of generation supplied by distributed energy resources At the end of five years, what percentage of power on your grid do you expect to be generated by DERs?	76%–100%	1.875	2.00	0.222
	51%–75%	1.550	1.50	0.295
	26%–50%	1.275	1.05	0.431
	0%–25%	0.675	0.80	0.358
Number of operating entities that will be contributing to generation How many entities (including the utility) will be operating generation sources on your grid in five years?	16+	1.650	1.75	0.387
	8–15	1.450	1.35	0.359
	4–7	1.325	1.15	0.406
	1–3	0.850	1.00	0.510
	0	1.025	1.00	0.690
Degrees of centrality Please choose the option that best describes your grid in five years.	One operating entity (probably the utility) issues control signals that the other entities must follow regarding the dispatch of resources and system regulation. The control signals specify what must be done and how to do it. For instance, "curtail generation by a specific amount by curtailing a specific resource."	1.100	1.15	0.692

Question	Answer	Mean	Median	Standard Deviation
	One operating entity (probably the utility) issues control signals that the other entities must follow regarding dispatch of resources and system regulation. The control signals specify what must be done but not how to do it. For instance, “curtail generation by a specific amount by whatever means are available.”	0.925	1.10	0.510
	One operating entity (probably the utility) issues control signals, but the other entities decide whether to comply.	1.625	1.60	0.265
	None of the operating entities issue control signals.	1.925	1.95	0.085
Sharing operational data between operating entities Please choose the option below that best describes your grid in five years.	Each operating entity shares operational data with every other operating entity.	0.650	0.55	0.583
	Each operating entity shares operational data with only one other operating entity, which provides general grid situational awareness to all operating entities (without widely sharing specific data).	1.050	1.25	0.609
	Operating entities share no operational data.	2.000	2.00	0
Operational focus Please choose the option that best describes your grid in five years.	Each operating entity will be able to determine its own objective.	1.550	1.60	0.373
	Each operating entity will be allowed some leeway in determining its own objective but will have to conform to certain parameters intended to ensure reliability.	0.725	0.75	0.406
Security alerts—content Please choose the option below that best describes your grid in five years.	Operating entities will not share information about cyber events.	1.750	2.00	0.444
	All operating entities will share ONLY statistics about the number and types of cyber events seen on their system.	1.150	1.30	0.420
	All operating entities will share statistics about the number and types of cyber events seen on their systems PLUS details of investigations into cyber events seen on their systems.	0.575	0.50	0.461
Security alerts—recipient In five years, who will receive the shared	Information about cyber events will not be shared.	1.750	2.00	0.444
	Information about cyber events will be shared with just the utility.	1.300	1.35	0.218

Question	Answer	Mean	Median	Standard Deviation
information concerning cyber events?	Information about cyber events will be shared with all operating entities.	0.775	0.95	0.406
Cloud products and services In five years, will operating entities on your grid be required to include cloud products or services in security and risk assessments?	No, cloud products and services will be excluded from security and risk assessments.	1.650	1.60	0.235
	Yes, all operating entities on my grid will be required to include cloud products and services in security and risk assessments.	0.750	0.65	0.495
Legacy vs. modern communication protocols In the next five years, will all operating entities on your grid be required to retire or transition away from legacy protocols?	No, operating entities on my grid will employ legacy protocols (e.g., Modbus, outdated versions of DNP3).	1.950	2.00	0.089
	Yes, all operating entities will be required to use modern, securely configured protocols (e.g., IEEE 2030.5 or OpenADR 2.0).	0.450	0.35	0.387
Wireless communication mediums In five years, what percentage of system-wide communications will be wireless?	21%+	1.250	1.25	0.573
	11%–20%	1.275	1.15	0.467
	0%–10%	0.950	0.85	0.720
Communication architecture and network segmentation Which option best describes your grid in five years?	No network segmentation procedures will be implemented.	1.750	1.90	0.336
	There will be some network segmentation procedure done on an *ad hoc* basis.	1.100	1.20	0.607
	Network segmentation procedures will follow state-level or national-level requirements or guidance including segmentation of IT, operational technology (OT), and business networks.	0.600	0.65	0.417
Including cybersecurity through interconnection agreements How do current interconnection agreements for your system address cybersecurity?	Cybersecurity is not mentioned in interconnection agreements.	1.875	2.00	0.222
	Cybersecurity is mentioned in interconnection agreements, but only in general terms.	1.450	1.40	0.373
	Cybersecurity is mentioned, and the interconnection agreements include specific cybersecurity requirements.	0.575	0.60	0.438

Question	Answer	Mean	Median	Standard Deviation
Monitoring of cyber threat intelligence Which will apply to ALL operating entities on your grid (including the utility) in the next five years?	One or more sources of cyber threat intelligence will be identified as being applicable to your grid.	0.775	1.00	0.400
	The sources identified above will be monitored, with security alerts issued to all operating entities as appropriate.	0.700	0.75	0.441
	The operating entities (including the utility) will create and exercise procedures for responding to credible and applicable cyber threat intelligence.	0.600	0.65	0.417
Application firewalls between entities Will application firewalls monitor communications between operating entities (including the utility) in the next five years?	No, application firewalls will not monitor communications between operating entities.	1.900	2.00	0.178
	Yes, application firewalls will monitor communications between operating entities (including the utility).	0.575	0.60	0.438
Supply chain and procurement Which of the following will be true of ALL operating entities on your grid (including the utility) in the next five years?	Supply chain cybersecurity risk management plan(s) (similar to what is described for BES in CIP-013-1) will be developed and implemented.	0.725	0.90	0.380
	Requests for proposals sent to vendors will include questions pertaining to cybersecurity and their product development processes.	0.675	0.65	0.449
	After the vendor is selected, the procurement language issued to the vendor will include baseline cybersecurity procurement language covering items such as device configurations and deactivation of unnecessary services. The baseline is meant to ensure the product is delivered in a state that supports the cybersecurity needs of the utility or operating entity.	0.625	0.55	0.455
IEEE 2030.5 Is it likely that an IEEE 2030.5 requirement will apply throughout your grid (both operating entities and the utility) in the next five years?	No, use of IEEE 2030.5 will not be required.	1.675	1.85	0.419
	Yes, use of IEEE 2030.5 will be required.	0.600	0.65	0.348
UL 1741 SA	No, use of UL 1741 SA will not be required.	1.500	1.75	0.628

Question	Answer	Mean	Median	Standard Deviation
Is it likely that a UL 1741 SA requirement will apply throughout your grid (both operating entities and the utility) in the next five years?	Yes, use of UL 1741 SA will be required.	0.550	0.55	0.462
Testing and certification Is it likely that requirements for the inspection, testing, and certification of distribution equipment (e.g., DERs, power conditioning equipment, safety equipment, and meters and instrumentation) will apply throughout your grid (operating entities and the utility) in the next five years?	No, inspection, testing, and certification will not be required.	1.550	1.85	0.634
	Yes, inspection, testing, and certification will be required.	0.525	0.50	0.438

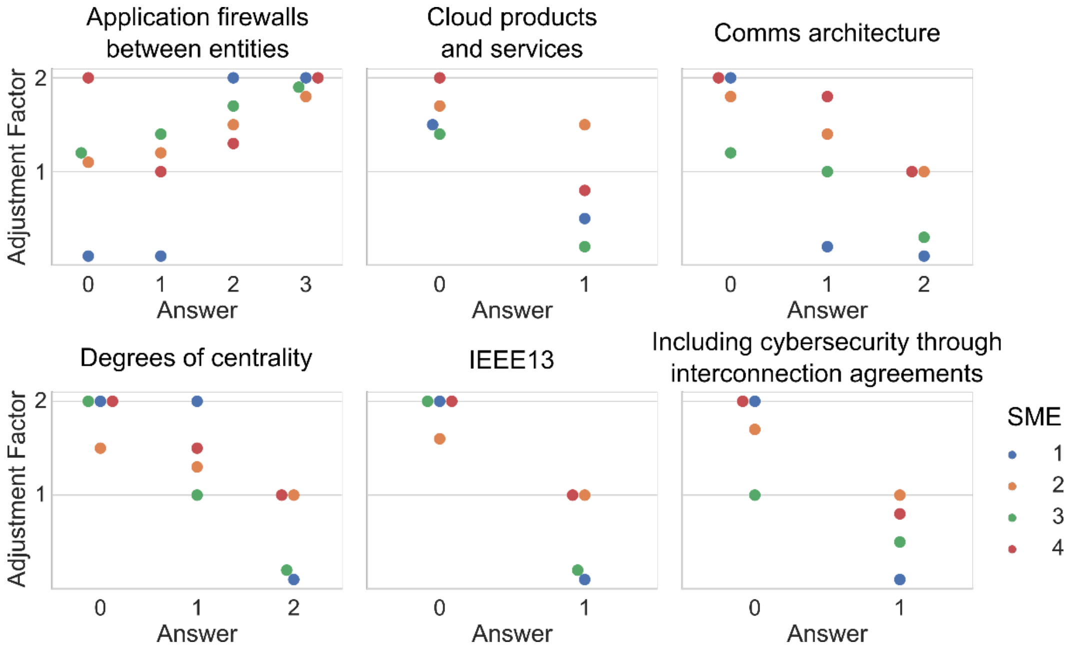

Figure C-4. Individual SME responses for the adjustment factors for six selected questions

Appendix D: Cyber100 Notional Use Case

The following hypothetical use case provides an example of how an organization could use Cyber100 Compass to model cybersecurity risks as it integrates large amounts of distributed energy resources. This use case is not intended to be comprehensive nor cover every event or condition that an organization might encounter. Rather, it is meant to serve as an illustrative example of a single instance of the Cyber100 Compass use.

The Situation

System A is a large, investor-owned utility serving more than 1 million customers and operating over a large geographic region in the state. The system has ambitious renewable energy transition targets (80%–90% renewable electricity and more than 65 GW of annual renewable energy capacity by 2050), with a goal of achieving 40% overall renewable energy and an additional 30 GW of renewable energy capacity deployed by 2030. System A seeks to add numerous new options, including the following, some of which the utility will own and operate, but most will be owned and operated by third-party entities. System upgrades include:

- New large-scale solar and wind power facilities
- Electric vehicle charging stations.

The Problem

Given the large number of customers supported by System A, the board of directors is increasingly concerned about how their plans to restructure the grid into a system of systems (where the constituent systems are operating entities of the different renewable resources) will change the cyberattack surface of their future energy system.

The Solution

System A's board of directors have tasked an employee, Gary, to use a new tool from the National Renewable Energy Laboratory (NREL)—the Cyber100 Compass—to understand and assess the cyber risks and mitigation strategies for their evolving energy system. Using the Cyber100 Compass will enable System A staff to quantify their cybersecurity risks, and it will allow Gary to present those risks to the System A Board members, executive leadership, and external stakeholders so that decisions and corrective actions can be made prior to the transition.

Risk Tolerance

The utility has already evaluated the risk of cyberattacks associated with System A's information technology (IT) systems. The IT networks are regularly hit with phishing, ransomware, and similar attacks, and the utility has a good sense of how much System A spends on average each year to recover from these IT events. As the utility transitions to more distributed energy resources and increases automation on their energy system, however, the utility suspects that their system is increasingly vulnerable to cyberattacks that could cause physical impacts, such as outages, harm to equipment, harm to employees, harm to the community, and denial of communications. These types of cyberattacks have not yet entered the utility's risk calculations.

Gary starts with the utility leadership (the board of directors and the senior executives). Gary and the chief risk officer begin a conversation with the operational technology (OT) security team, system planners, and other senior executives about risks from cyberattacks that cause physical impacts. The goal of this conversation is to determine which total annual costs are acceptable across a spectrum of cyberattack probabilities. As total costs increase, the acceptable probability of experiencing those costs decreases.

To formalize the decision, System A presents all the stakeholders (including the senior executives) with a series of scenarios. Each scenario includes a question. For instance:

> **Scenario 1: More than $1,000**
>
> A cyberattack on OT systems results in physical impacts to the grid. The total cost to the organization, including lost revenue, recovery, etc., is more than $1,000.

In this scenario, what is the acceptable probability that a cyberattack costing more than $1,000 will occur in any given year? Please write the probability as a percentage between 0 and 100.

Scenario 2 is identical except that the cyberattack would cost System A more than $10,000. Scenario 3 introduces a cyberattack costing more than $100,000, and so on. The answers are summarized in Table D-1.

Table D-1. Example Risk Tolerance Inputs

Total Costs From Attacks	More than $1,000	More than $10,000	More than $100,000	More than $1,000,000	More than $10,000,000
Acceptable probability of cyberattacks costing this amount in any given year	95%	35%	15%	2%	0.1%

This enables System A to generate the risk tolerance curve shown in Figure D-1.

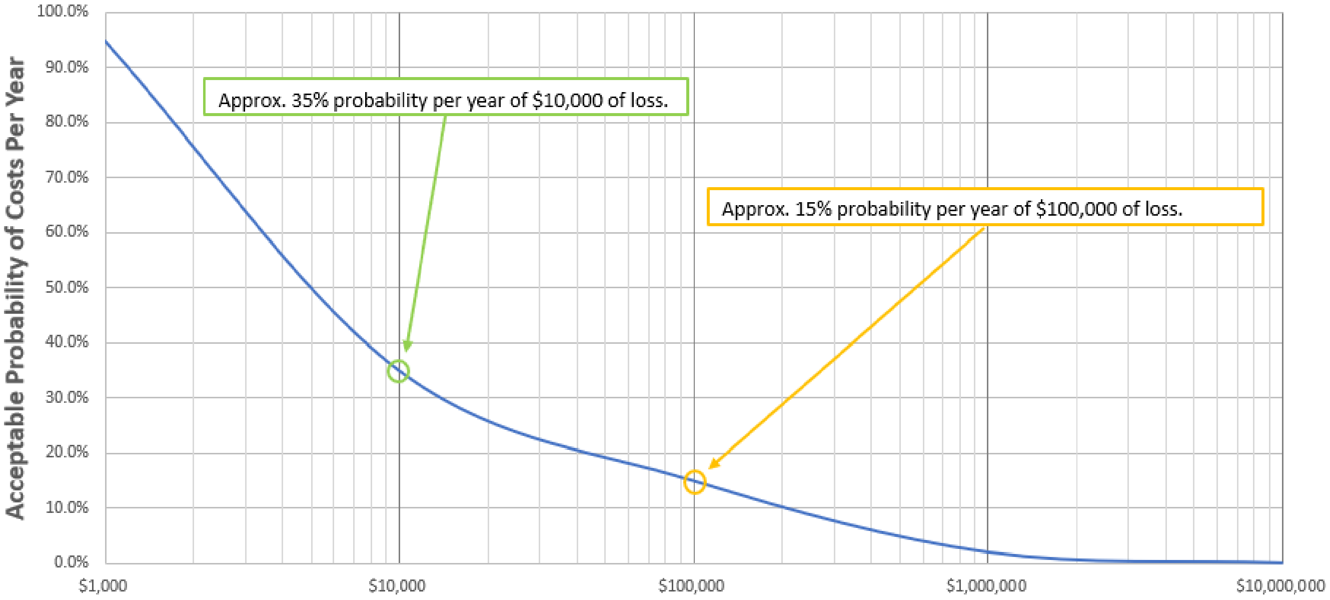

Figure D-1. Example risk tolerance curve

The risk tolerance curve provides a probabilistic risk assessment that helps organizations quantify risk. For more information on risk tolerance curves and probabilistic risk assessment, see Chapter 3 of *How to Measure Anything in Cybersecurity Risk* (Hubbard and Seiersen 2016, 35).

The Cyber100 Compass will use the values in this section (together with the input values in the next section) to create System A's recommended expectation-of-loss curve.

Events

The next step is for System A to provide inputs for the events section. These inputs allow users to tell the Cyber100 Compass the value placed on avoiding certain types of events that could be caused by a successful cyberattack. In the resilience space, these values are often called "avoided costs." System A is asked to provide the maximum avoided costs for events given certain constraints that define three levels of impact: low, moderate, and high.

The sources of information needed to estimate event avoidance values depend on each event category in the Cyber100 Compass. Gary uses a few publicly available tools to help estimate these costs, and then he validates these estimates with System A's finance and risk management committee. The event categories that require Gary to create event avoidance values in the proof-of-concept application are:

- Power outage
- Harm to equipment
- Harm to employees (of the utility or operating entity)
- Harm to the community (e.g., overloading equipment to create a wildfire)
- Loss of productivity or efficiency arising from loss of communications.

Power Outage

Gary uses a free, public tool, called the Interruption Cost Estimate (ICE) Calculator,[2] developed by Lawrence Berkely National Laboratory, Resource Innovations, and sponsoring utilities, which helps System A estimate the costs of power interruptions. The ICE Calculator uses common reliability metrics: the System Average Interruption Frequency Index (SAIFI), or how many minutes of electric interruptions that an average customer experienced per year; the System Average Interruption Duration Index (SAIDI), or the number of times an electrical interruption occurred per year; and the Customer Average Interruption Duration Index (CAIDI), or the average time it took to restore power after an electrical interruption. The metrics are already available to Gary because System A already collects these metrics for their own reliability tracking and reporting. Using the ICE Calculator, Gary can estimate the costs of a power outage and estimate the low, moderate, and high impacts of a power outage to customers. Risk managers then review these estimates and decide on values based on the level of impact.

[2] See https://icecalculator.com/home.

Harm to Equipment

Again, Gary turns to a free, public tool, called the Customer Damage Function Calculator (CDF),[3] developed by NREL. Estimating the harm to equipment will require some additional conversations with security and operations teams to determine the kinds of equipment that could require repair or replacement after a damaging cyberattack. This can include damage or replacement costs to supervisory control and data acquisition system equipment, such as remote terminal units, programmable logic controllers, or human machine interfaces; or possibly repair or replacement costs for resources such as solar photovoltaics, wind turbine components, and electric vehicle chargers. Gary works with operations personnel to estimate the average costs of the kinds of equipment that could be damaged in a cyberattack. Then, estimating low, moderate, and high impacts, Gary uses the CDF Calculator to estimate damage costs, which uses the equation CxNxP, where C = the average cost of equipment repair or replacement, N = the number of pieces of equipment damaged, and P = the probability of damage from an outage. This could help Gary and risk managers estimate the equipment costs after a cyberattack at different levels of impact.

Harm to Employees

The U.S. Department of Labor developed the "OSHA $afety Pays" Estimator to help organizations estimate the impact of occupational injuries and profitability losses based on the type of injury, the profit margins of the organization, and the number of employees injured.[4] Gary can estimate the cost and impact level of employee harm based on a potential type of injury from a cyberattack—for instance, a burn could be less costly than an electric shock. Further, Gary can also estimate a worst-case scenario, where high impact means the death of an employee. Other sectors and agencies such as the U.S. Environmental Protection Agency and the U.S. Department of Transportation, use a common mortality risks valuation process called the "value per statistical life" to estimate the costs System A would be willing to pay to reduce the probability of an employee's death. Once Gary estimates a range of cost estimates for low, moderate, and high impact, Gary works with human resources personnel to validate these estimates.

Harm to Community

Gary considers other power systems events that have caused harm to the community and decides to estimate the costs of a cyberattack creating power flow violations that start a forest fire. Although this is likely to be a low probability event, the impact could be very high. Gary can use past utility-caused wildfire liability cases to estimate costs based on the number of System A customers and their geographic locations and create hypothetical scenarios of low, moderate, and high impact. These estimates are discussed and validated by risk management personnel.

Loss of Productivity or Efficiency Arising from Loss of Communications

To estimate the impact from a loss of communications, Gary can use the CDF Calculator, which allows him to estimate downtime costs—calculated as the cost per hour if some or all of System A's operations are idle, multiplied by the hours of downtime. He can also estimate process

[3] See https://cdfc.nrel.gov/.
[4] See https://www.osha.gov/safetypays/estimator.

interruption and restart costs, calculated by multiplying the average hourly employee costs (wage plus overhead) by the hours of staff time needed to reset or reestablish communications with assets. Gary can use the calculator to estimate these costs at different impact levels. Both risk management and human resources personnel are helpful to validate these estimates.

Event Avoidance Values

At each stage in the process of creating event avoidance values, Gary works with a team of System A stakeholders to validate that these estimates are reasonably consistent with how finance and risk managers would evaluate these costs. By identifying tools to help him estimate these costs, Gary helps System A leadership identify the maximum avoided costs of the financial impacts that could occur based on the description of the cyber events and their levels of impact as described in the Cyber100 Compass.

Conditions

The next step is for System A to select conditions that it expects to apply to its energy systems in the future and which will impact cybersecurity.

For the Cyber100 Compass, the conditions describe any constraints, resources, requirements, controls, or other factors that modify the cybersecurity risks of System A's energy transformation plans. To gather the necessary input information for Cyber100 Compass, Gary must meet with internal System A stakeholders as well as external stakeholders.

For example, to gather information related to the regulatory environment, Gary meets with System A's government relations office to understand state-level resources that could be valuable to System A in case of a cyberattack. Several policy practices that are implemented in other states could come to System A's state in the future. Gary discusses these future possibilities with the corporate governance team, which can then leverage their relationships with state legislators to discuss any potential critical infrastructure cybersecurity policies that might occur within the next 5 years.

Running Cyber100 Compass

Gary has now consulted with the appropriate stakeholders, gathered information on System A's current and future conditions, and input this information. Depending on the number of Monte Carlo simulations selected, running the Cyber100 Compass takes at most a few minutes. The Cyber100 Compass generates a report and visualizations that capture the risk calculations made based on Gary's inputs. Gary collects the reports and creates a cover page summarizing the findings for senior executive decision makers. The Cyber100 Compass can be run multiple times to test various future scenarios and to adjust inputs as conversations with stakeholders evolve.

Conclusions

The Cyber100 Compass has given Gary and System A an actionable and quantifiable risk assessment for their future energy system. System A can now make more informed and better decisions to ensure that the next stage of its energy system's evolution is designed with security in mind.

System A expects to use the Cyber100 Compass again in approximately 5 years to quantify the risks involved with the following stage of upgrades (the next 5-year increment). At that time, System A expects that a new version of the Cyber100 Compass will be available to provide even better risk insights.

Appendix E: Condition Inputs

The following text was used for the condition inputs for the Cyber100 Compass proof-of-concept application.

Condition questions allow users to tell Compass things they believe to be true about their current and future grid that will impact cybersecurity risk. Conditions in Compass are similar to what NIST defines as a predisposing condition:

> A condition that exists within an organization, a mission/business process, enterprise architecture, or information system including its environment of operation, which contributes to (i.e., increases or decreases) the likelihood that one or more threat events, once initiated, will result in undesirable consequences or adverse impact to organizational operations and assets, individuals, other organizations, or the Nation.

Compass extends this idea from the present into the future; many of the **Condition** questions ask users what they expect to be true about their grid in the next five years.

Condition questions are meant to elicit information; they are not meant to be—and should not be interpreted to be—policy recommendations. Further, Compass is not primarily focused on physical controls for security. The purpose of Compass is to assess the risk of cyberattacks producing physical impacts on energy systems, not physical security breaches that have cybersecurity impacts.

Condition questions are spread across the following five categories, each with a corresponding section:

- Changes to Grid Topology
- Changes to System-of-Systems Architecture
- Communications
- Security Controls
- Regulatory Environment.

Users can expect to complete the **Condition** questions in approximately two hours, although exact completion times will vary. It may be helpful to refer to the following information sources as you progress through the questions:

- Utility personnel: chief security officer, chief technology officer, chief information officer, system planners, control engineers, power engineering teams, communication engineers, chief financial officer, corporate governance team, government affairs office
- System planning data: design documents, power purchase agreements, capacity expansion models, projections of load growth
- Persons involved in the development and implementation of interconnection agreements
- Current or future third-party operating entities
- State public utility commissions (PUCs)
- State, regional, or national service organizations (e.g., the National Rural Electric Cooperative Association or one of its state-level associations)

Defining Terms

- **Your grid** refers only to the specific electricity grid for which you are trying to assess future risks. (You are not being asked, for instance, to provide answers that apply to the national U.S. grid.)
- **Operating entity** refers to a grid participant that supplies a significant contribution to generation, storage, or load control. Operating entities may include, for instance, operators of wind or solar farms, battery banks, or demand response capabilities. Solar power aggregators would also be considered a type of operating entity. The utility for a grid service area is also an operating entity; although, in the **Condition** questions the role of the utility is often called out for clarity.
- **Distributed energy resources (DERs)** are, for the purposes of Compass, small-scale energy generation and storage technologies producing less than 10 megawatts (MW) of power.
- **Cyber events** (or simply **events**) are any observable instance of an attempted or successful cyberattack. For simplicity, Compass does not distinguish between what cybersecurity literature calls cyber events (any observable occurrence of a possible or attempted cyberattack) and cyber incidents (one or more cyber events that succeed in having a negative impact on the target system). Compass uses **cyber event** to cover both concepts. Compass focuses on cyber events that have physical impacts on operational technology (OT) networks.

How Compass Adjusts Risk Based on Conditions

Compass begins with a baseline risk for each type of cyber event, then adjusts this risk based on the **Condition** questions answered by the user. Each of your answers refines the risk estimate by providing information specific to the system under consideration. These refinements improve the probabilistic calculations that quantify the cybersecurity risks for your grid.

Future Projections

Most questions ask about expected conditions on the grid within a five-year period. When interpreting the output after Compass is run, users should likewise consider the annual loss expectancy to be a projection no more than five years into the future.

Uncertainty in Answers

Because Compass users are asked about future aspects of their grids, there is some amount of uncertainty. Users of Compass will not have perfect knowledge regarding their grids in the five-year timespan identified in most questions. If a question asks for one and only one answer, Compass gives users the opportunity to select "**unsure**." However, users are encouraged to use this option as little as possible. Selecting more decisive answers—based on your business plans, information from third parties, trends, your own foresight, and other sources—produces results more specific to your own grid. The more questions decisively answered, the better Compass can estimate risk and calculate recommended expectation of loss values for cybersecurity events.

Question Types

"One and only one" questions require users to input only one answer. Users are required to select an answer to every "one and only one" question before Compass will run, even if the answer is "**unsure**."

"All that apply" questions allow users to select multiple answers from a list. For these questions, users have the option to indicate if none of the answers apply.

Headings

Text for each of the **Condition** questions includes the following headings:

- **Background** provides necessary context for each question.
- **Assumption(s)** gives a qualitative explanation of how the answer will affect the risk estimate. The phrase **"Compass adjusts risk downward"** means that Compass assumes the answer will make negative impacts less likely (a good thing). **"Compass adjusts risk upward"** means that Compass assumes the answer will make these negative impacts more likely (a bad thing). Note that selecting "**unsure"** means your answer will not adjust risk in either direction, essentially eliminating that question from subsequent risk calculations.
- **Question** is the actual request for information. It is followed by one or more answer options. Please see above section "Question Types" for more information on the format of questions.
- **References** provide citations and additional background information where needed. (Note: Not all questions include references.)

Changes to Grid Topology

The transition to high levels of renewable energy will require rethinking how generation is owned and managed. It is generally assumed that the inclusion of high levels of renewables will make the grid more distributed, but the nature of this distribution impacts grid security. Also, this raises the question of grid topology—how the distributed elements of the grid are arranged and interconnected. The questions below solicit information on this topic.

Sources of information useful in filling out the questions in this section could come from design documents, power purchase agreements, capacity expansion models, projections of load growth, and other system planning data available to system engineers, utility boards of directors, or other executive leadership.

Percentage of Generation Not Operated by the Utility

Background

How much of the power used by the future grid will be generated from sources not operated by the utility? This power may come from power purchase agreements with operators of renewable generation, small grid-connected installations (such as rooftop solar), or other sources.

Assumption

Compass adjusts risk upward for more generation not operated by the utility. While non-utility operators may be quite secure, the utility cannot be entirely certain of this. In the options below, risk increases as the range of percentages increases.

Question

At the end of five years, what percentage of power on your grid do you expect to be generated by sources not operated by the utility? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- 76%–100%
- 51%–75%
- 26%–50%
- 0%–25%
- Unsure.

Largest Non-Utility Operating Entity

Background

If a single non-utility operating entity were taken offline due to cyberattack, the impact of this event would depend on the percentage of power supplied by that operating entity. (Also, please consider energy capacity if you expect your system's dispatchable resources to be dominated by storage.)

Assumption

Compass adjusts risk upward if a single non-utility operating entity supplies a large percentage of total generation. While non-utility operators may be quite secure, the utility cannot be entirely certain of this. In the options below, risk increases as the range of percentages increases.

Question

At the end of five years, what percentage of power capacity (or in the case of storage, energy capacity) on your grid will be supplied by the largest non-utility operating entity? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- 76%–100%
- 51%–75%
- 26%–50%
- 0%–25%
- Unsure.

Percentage of Generation Supplied by Distributed Energy Resources

Background

Future generation may come from many distributed energy resources (DERs), defined here as resources producing 10 MW or less.

Assumption

Compass adjusts risk upward if more power is generated by many small installations, which would tend to increase the attack surface of the grid. In the options below, risk increases as the range of percentages increases.

Question

At the end of five years, what percentage of power on your grid do you expect to be generated by DERs? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- 76%–100%
- 51%–75%
- 26%–50%
- 0%–25%
- Unsure.

Number of Operating Entities That Will Be Contributing to Generation

Background

The generation not operated by the utility will be operated by one or more other entities. How many of these distinct non-utility operating entities will be participating?

Assumption

Compass adjusts risk upward as the number of operating entities increases. An increase in the number of operating entities results in an increase in the attack surface of the grid.

Question

How many entities (including the utility) will be operating generation sources on your grid in five years? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- 16+
- 8–15
- 4–7
- 1–3
- 0
- Unsure.

Changes to System-of-Systems Architecture

Future high-renewable grids may function as collections of operating entities, each supplying some percentage of the overall generation mix. This raises the question of how these operating entities will coordinate and dispatch resources at a local level to ensure operational stability and reliability. The questions below describe qualities of system-of-systems architecture that might be used to integrate operating entities on future grids.

Sources of information useful in filling out the questions in this section could come from the utility board of directors, chief security officer, chief technology officer, chief information officer, system planners, control engineers, and power engineering teams.

Degrees of Centrality

Background

Systems of systems can be categorized by the degree to which they are "centralized" (i.e., the degree to which they operate with a central entity establishing priorities and issuing control signals). While centralization has some advantages, it can also create a single point of failure for the system of systems. For more information, see the Department of Defense's *Systems Engineering Guide for Systems of Systems*.

Assumption

Compass adjusts risk upward as the system-of-systems architecture becomes more centralized. An increase in centrality means that if the controlling entity is compromised through cyberattack, the attackers have a better chance to issue malicious controls to other entities.

Question

Please choose the option below that best describes your grid in five years. The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- One operating entity (probably the utility) issues control signals that the other entities must follow regarding dispatch of resources and system regulation. The control signals specify what must be done and how to do it. For instance, "curtail generation by a specific amount by curtailing a specific resource."
- One operating entity (probably the utility) issues control signals that the other entities must follow regarding dispatch of resources and system regulation. The control signals specify what must be done but not how to do it. For instance, "curtail generation by a specific amount by whatever means are available."
- One operating entity (probably the utility) issues control signals, but the other entities decide whether to comply.
- None of the operating entities issue control signals. All operating entities make decisions based on their own objectives and market signals.
- Unsure

Sharing Operational Data Between Operating Entities

Background

Many types of data can be shared to better coordinate constituent systems within a system-of-systems environment. This includes individual system development plans and funding profiles.

However, the sharing of such data can be a contentious issue, as there may be business or institutional reasons not to share. Sharing operational data can be a double-edged sword: it allows operating entities to better coordinate system-wide performance but may allow attackers opportunities to compromise operational data in transit. For more information, see the Department of Defense's *Systems Engineering Guide for Systems of Systems*.

Assumption

Compass adjusts risk upward as more operational data is shared between operating entities. Although such sharing may improve system-wide performance, it also introduces opportunities for data-spoofing attacks and other attacks on data in transit.

Question

Please choose the option below that best describes your grid in five years. The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- Each operating entity shares operational data with every other operating entity.
- Each operating entity shares operational data with only one other operating entity, which provides general grid situational awareness to all operating entities (without widely sharing specific data).
- Operating entities share no operational data.
- Unsure

Operational Focus

Background

In some systems of systems, all constituent systems (e.g., operating entities) share the same objective, and adherence to that objective is rigorously enforced. In other systems of systems, the constituent systems may have differing objectives, and the assumption is that the actions of all the systems will converge on a desired outcome. For instance, the utility may have grid stability and reliability as its primary objectives, while another operating entity (such as a third-party solar operator) may prioritize return on investment. For more information, see the Department of Defense's *Systems Engineering Guide for Systems of Systems*.

Assumption

Compass adjusts risk downward with increasing alignment between the objectives of the operating entities. This is because cyber attackers might be able to leverage differences among the objective functions of the operating entities to create grid instability.

Question

Please choose the option below that best describes your grid in five years. The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- Each operating entity will be able to determine its own objective.

- Each operating entity will be allowed some leeway in determining its own objective but will have to conform to certain parameters intended to ensure reliability.
- Unsure

Security Alerts – Content

Background

Operating entities will inevitably experience cybersecurity events, ranging from simple probes of perimeter defenses up to successful breaches that disrupt operations, damage equipment, or endanger life. What information about these events will the operating entities share with the utility and with each other?

Assumption

Compass adjusts risk downward with increased information sharing. Information sharing gives operating entities better situational awareness regarding the threat landscape and advanced notice of possible vectors and techniques for cyberattacks.

Question

Please choose the option below that best describes your grid in five years. The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- Operating entities will not share information about cyber events.
- All operating entities will share ONLY statistics about the number and types of cyber events seen on their system.
- All operating entities will share statistics about the number and types of cyber events seen on their systems PLUS details of investigations into cyber events seen on their systems.
- Unsure

Security Alerts—Recipients

Background

When operating entities share information about cyber events, with whom will they share it?

Assumption

Compass adjusts risk downward when information is distributed to more parties. A wider distribution enables more entities to act on the information that is shared.

Question

In five years, who will receive the shared information concerning cyber events? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- Information about cyber events will not be shared.
- Information about cyber events will be shared with just the utility.

- Information about cyber events will be shared with all operating entities.
- Unsure

Cloud Products and Services

Background

The future grid will likely include some form of cloud services to take advantage of benefits in redundancy, flexibility, reliability, and uptime. For information about benefits, see Sandia National Laboratories' *Roadmap for Photovoltaic Cyber Security*. However, cloud migration raises security concerns shared between energy systems and cloud service providers. For more information on cloud security, see the GSA's Cloud Information Center.

Assumption

Compass adjusts risk downward for systems that conduct a security assessment of cloud products or services, including documented impact level and the cloud service model and applicable security controls per NIST SP 800-53 Rev. 5: Security and Privacy Controls for Information Systems and Organizations.

Question

In five years, will operating entities on your grid be required to include cloud products or services in security and risk assessments? The requirement may be included in interconnection agreements or communicated to operating entities by some other means. The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- No, cloud products and services will be excluded from security and risk assessments.
- Yes, all operating entities on my grid will be required to include cloud products and services in security and risk assessments.
- Unsure

Communications

The future grid will likely be composed of a heterogeneous mix of communications protocols, physical mediums, and network architectures for system management and control. This raises questions related to the security of communications as data moves between different networking communication layers. As identified in the Power Systems book series chapter, Information and Communication Infrastructures in Modern Wide-Area Systems, the questions below address communication security in three critical areas: **protocols**, or the way data are formatted; **physical mediums** by which data are transmitted; and **architecture** of communication systems.

Sources of information useful in filling out the questions in this section could come from the utility chief security officer, chief technology officer, chief information officer, communication engineers, system planners, and others involved in the development and implementation of interconnection agreements.

Legacy vs. Modern Communications Protocols

Background

Many utilities in North America today use communications protocols like Modbus and outdated versions of DNP3 that were not originally designed with security in mind. For the purposes of Compass, these protocols built without native security features are considered legacy protocols. Modern protocols, such as IEEE 2030.5 and OpenADR, were designed with at least some built-in encryption and authentication features.

Assumption

Compass adjusts risk upward for entities using legacy communications protocols (like Modbus and outdated versions of DNP3) without native cybersecurity features. Using modern protocols built with security in mind is still the most cost-effective and secure approach to providing encryption, authentication, and authorization.

Question

In the next five years, will all operating entities on your grid be required to retire or transition away from legacy protocols? Legacy protocols include protocols like Modbus or outdated versions of DNP3. More modern and secure protocols include IEEE 2030.5 or OpenADR 2.0 with implemented trust and cryptography features. The requirement may be included in interconnection agreements, in standards the utility or operating entity chooses to apply, or communicated to operating entities by some other means. The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- No, operating entities on my grid will employ legacy protocols (e.g., Modbus, outdated versions of DNP3).
- Yes, all operating entities will be required to use modern, securely configured protocols (e.g., IEEE 2030.5 or OpenADR 2.0).
- Unsure

References

A discussion of communications protocols and security recommendations can be found in:

- *Cyber Security Primer for DER Vendors, Aggregators, and Grid Operators*, Sandia National Laboratories Report (2017)
- A Survey of Protocol-Level Challenges and Solutions for Distributed Energy Resource Cyber-Physical Security, *Energies* (2018)
- *Recommendations for Data-in-Transit Requirements for Securing DER Communications*, Sandia National Laboratories Report (2020)
- *The Utility Regulator's Role in Promoting Cybersecurity: Resilience, Risk Assessment, and Standards*, USAID-NARUC Publication (2020)

Wireless Communication Mediums

Background

The physical means by which information flows between geographically dispersed energy systems has shifted from predominantly wired communications like copper lines and fiber optic cables towards wireless communications like cellular, microwave, and satellite taking on an increasingly greater share of the communication mix.

Assumption

Compass adjusts risk upward for entities in which wireless communications is heavily used. Wireless connections can be misconfigured or poorly implemented, potentially allowing adversaries to gain access to high-level functions, disrupt operations, or launch other malicious activity like denial-of-service attacks.

Question

In five years, what percentage of system-wide communications will be wireless? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- 21%+
- 11%–20%
- 0%–10%
- Unsure.

References

A discussion of cybersecurity challenges and mitigations in wireless communications can be found in:

- Analysis and mitigation of vulnerabilities in short-range wireless communications for industrial control systems, *International Journal of Critical Infrastructure Protection* (2012)
- Information and Communication Infrastructures in Modern Wide-Area Systems, in: *Wide Area Power Systems Stability, Protection, and Security*, Power Systems, Springer, Cham (2012)

Communications Architecture and Network Segmentation

Background

The communications architecture of the future grid is likely to be highly heterogeneous, with the central controllers exchanging information with onsite and remote field devices, as well as external vendors. Network segmentation can restrict unauthorized access to critical parts of the system and limit the ability to pivot between different network segments.

Assumption

Compass adjusts risk downward for entities that have an organizational policy for network segmentation. According to NIST SP 800-82 Rev. 2: Guide to Industrial Control Systems (ICS) Security, network segmentation is one of the most effective architectural approaches to significantly reducing access to sensitive information from malicious actors or human error.

Question

Which option below best describes your grid in five years? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- No network segmentation procedures will be implemented.
- There will be some network segmentation procedures done on an ad-hoc basis.
- Network segmentation procedures will follow state-level or national-level requirements or guidance including segmentation of IT, operational technology (OT), and business networks.
- Unsure.

References

For an overview of network segmentation best practices, see *Network Perimeter Defense: Best Practices in Network Segmentation*, EnergySec Publication (2014).

For a description of NIST recommendations on network segmentation for industrial control systems, see NIST SP 800-82 Rev. 2: Guide to Industrial Control Systems (ICS) Security (2015).

Security Controls

This section lists only security controls applied at the system-of-systems level. These controls fall into three basic categories: 1) requirements for cybersecurity controls internal to an operating entity, 2) cybersecurity controls for information passing between operating entities, and 3) cybersecurity controls for the supply chain. These controls might be written into interconnection agreements for operating entities or introduced through some other contracting mechanism.

Sources of information useful in filling out the questions in this section could come from the utility chief security officer, chief technology officer, chief information officer, chief financial officer, those involved in the development and implementation of interconnection agreements, and current or future third-party operating entities.

Including Cybersecurity Through Interconnection Agreements

Background

The inclusion of cybersecurity in interconnection agreements is a relatively new development. Interconnection agreements or standards are formal documents dictating how certain energy systems including DERs and larger controllable loads (like electric vehicle chargers) can legally connect to the electricity grid. Some agreements mention cyber only in very general terms, with no specific requirements enumerated, while others are more specific. One example comes from the State of Minnesota's Technical Interconnection and Interoperability Requirements, which require DER operators to consider physical and front panel security, network security, and communication interface security.

Assumption

Compass adjusts risk downward if cybersecurity is mentioned in interconnection agreements, and more so if specific requirements are included. Any mention of cyber indicates a concern that will hopefully translate to action, with more specific language indicating even more concern. If cybersecurity is not mentioned at all, risk is adjusted upward.

Question

How do current interconnection agreements for your system address cybersecurity? Answer based on interconnection agreements in place today or in the near future (and which will likely still be in effect in five years). The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- Cybersecurity is not mentioned in interconnection agreements.
- Cybersecurity is mentioned in interconnection agreements, but only in general terms.
- Cybersecurity is mentioned, and the interconnection agreements include specific cybersecurity requirements.
- Unsure.

Monitoring of Cyber Threat Intelligence

Background

Cyber threat intelligence is information collected about potential threats. It may be industry-specific (e.g., focus on threats to utilities) or general (i.e., covering many industries).

Assumption

Compass adjusts risk downward if utilities and operating entities monitor cyber threat intelligence and have procedures in place to act on intelligence that represents a credible threat.

Question

Which of the following will apply to ALL operating entities on your grid (including the utility) in the next five years? The efforts listed below may be written as requirements in the interconnection agreement or communicated to operating entities as a requirement by some other means. Please select all that apply. If none are applicable, select "None apply."

- One or more sources of cyber threat intelligence will be identified as being applicable to your grid.
- The sources identified above will be monitored, with security alerts issued to all operating entities as appropriate.
- The operating entities (including the utility) will create and exercise procedures for responding to credible and applicable cyber threat intelligence.
- None apply.

Application Firewalls Between Entities

Background

Many organizations are interconnected to a much greater extent than in the past, with partner organizations sharing persistent connections to share data and control signals. One attack path that has proven effective involves compromising one partner and pivoting to others. High-renewable grids that function as a collection of connected operating entities could be targeted with such an attack. One possible mitigation would be to set up application firewalls at the communication touchpoints between the operating entities. These firewalls (working at the application layer of the Open Systems Interconnection model) could detect and/or stop malicious signals passing between operating entities and the utility.

Assumption

Compass adjusts risk downward if utilities and operating entities apply application firewalls to monitor inter-entity traffic.

Question

Will application firewalls monitor communications between operating entities (including the utility) in the next five years? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- No, application firewalls will not monitor communications between operating entities.
- Yes, application firewalls will monitor communications between operating entities (including the utility).
- Unsure.

Supply Chain and Procurement

Background

Cybersecurity at all levels of the grid depends on the security of devices and systems that constitute the grid. The relatively recent addition of supply chain risk management to NERC CIP (via CIP-013-1) illustrates the significance of this issue to bulk electric system (BES) operations. The same concern applies equally to the distribution system.

Assumption

Compass adjusts risk downward if more efforts are made to address cyber risk in supply chain and procurement. This is because these efforts can be a critical part of managing cyber risk.

Question

Which of the following will be true of ALL operating entities on your grid (including the utility) in the next five years? The efforts listed below may be written as requirements in the interconnection agreement or communicated to operating entities as a requirement by some other means. Please select all that apply. If none are applicable, select "None apply."

- Supply chain cybersecurity risk management plan(s) (similar to what is described for BES in CIP-013-1) will be developed and implemented.

- Requests for proposals sent to vendors will include questions pertaining to cybersecurity and their product development processes. (Questions can be adapted from vendor websites or other industries.) The vendors will be required to submit answers to these questions when they submit their proposals; the utility or operating entities will review these answers, and the information about vendor security will be factored into the vendor selection decision.
- After the vendor is selected, the procurement language issued to the vendor will include baseline cybersecurity procurement language covering items such as device configurations and deactivation of unnecessary services. The baseline is meant to ensure the product is delivered in a state that supports the cybersecurity needs of the utility or operating entity.
- None apply.

References

More about managing cyber supply chain risk can be found in *Cyber Supply Chain Risk Management for Utilities—Roadmap for Implementation*, Utilities Telecom Council Report (2015)

More about cybersecurity procurement language can be found in *Cybersecurity Procurement Language for Energy Delivery Systems*, 9th Annual Cyber and Information Security Research Conference (2014).

More about cybersecurity procurement language can be found in *Cyber Security Procurement Language for Control Systems*, U.S. Department of Homeland Security Report (2009).

Regulatory Environment

As grid architecture continues to evolve from its original design based on central generation and control, larger and larger portions of the distribution grid fall outside federal jurisdiction and are therefore not subject to NERC CIP. This is especially relevant for DERs, which are typically connected at the local electric distribution level and not to the sub-transmission or transmission network (see *GMLC Survey of Distributed Energy Resource Interconnection and Interoperability Standards*). States, PUCs, and other entities have begun trying to address distribution-level cybersecurity via regulatory mechanisms that necessarily affect the development of the future grid.

The questions in this section draw from examples already in place or under consideration in different U.S. states (as described in the “Background” text for each condition below). Your answers should indicate if you think similar regulatory mechanisms might be in play in your state over the next five years. If you are unsure how to answer a certain question, you may select the **“Unsure”** answer choice.

Sources of information useful in filling out the questions in this section could come from the state PUC; the utility’s corporate governance team; the utility’s government affairs office; or state, regional, or national service organizations (e.g., the National Rural Electric Cooperative Association or one of its state-level associations).

IEEE 2030.5

Background

California Rule 21 is an interconnection tariff requiring most new residential and commercial DER systems to communicate with the host utility using IEEE 2030.5, a smart energy standard that includes requirements for transport-level security and strong encryption.

Assumptions

Compass adjusts risk downward for systems requiring the use of IEEE 2030.5. Adherence to this standard helps secure communications to and from DERs, thus potentially lowering the risk of cyberattack at specific points within the system.

Question

Is it likely that an IEEE 2030.5 requirement will apply throughout your grid (both operating entities and the utility) in the next five years? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- No, use of IEEE 2030.5 will not be required.
- Yes, use of IEEE 2030.5 will be required.
- Unsure

References

A high-level overview of IEEE 2030.5 can be found in Securing California Rule 21 Networks, SunSpec Alliance Cybersecurity Webinar (2018).

More information about the current state of DERs and interconnection standards can be found in *GMLC Survey of Distributed Energy Resource Interconnection and Interoperability Standards,* NREL and Lawrence Berkeley National Laboratory Report (2020).

UL 1741 SA

Background

California and Hawaii require all advanced inverter functions to be certified to UL 1741 SA, a product safety standard that establishes manufacturing and testing requirements for smart inverters.

Assumptions

Compass adjusts risk downward for systems requiring the use of UL 1741 SA. Adherence to this standard helps ensure smart grid inverters meet rigorous requirements prior to interconnection.

Question

Is it likely that a UL 1741 SA requirement will apply throughout your grid (both operating entities and the utility) in the next five years? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- No, use of UL 1741 SA will not be required.
- Yes, use of UL 1741 SA will be required.
- Unsure.

References

More information about the current state of DERs and interconnection standards can be found in *GMLC Survey of Distributed Energy Resource Interconnection and Interoperability Standards*, NREL and Lawrence Berkeley National Laboratory Report (2020).

Testing and Certification

Background

Some states are beginning to consider broader testing and certification requirements for DER interconnection. For example, proposed changes to Michigan's Interconnection and Distributed Generation Standards rule set would add testing and certification requirements. Under this proposed rule change, if the interconnection application requires telecommunications, cybersecurity, data exchange, or remote controls operation, successful testing and certification of these items must be completed prior to deployment.

Assumptions

Compass adjusts risk downward for systems with testing and certification requirements. These requirements introduce rigor into the interconnection process, thus potentially lowering the risk of cyberattack at interconnection points for DERs within the system.

Question

Is it likely that requirements for the inspection, testing, and certification of distribution equipment (e.g., DERs, power conditioning equipment, safety equipment, and meters and instrumentation) will apply throughout your grid (operating entities and the utility) in the next five years? The answers are ordered from highest risk to lowest risk. Choose one and only one answer.

- No, inspection, testing, and certification will not be required.
- Yes, inspection, testing, and certification will be required.
- Unsure.

References

More information about cybersecurity regulations in U.S. states can be found in:

- *Improving the Cybersecurity of the Electric Distribution Grid Phase 1 Report: Identifying Obstacles and Presenting Best Practices for Enhanced Grid Security*, Vermont Law School Report (2019)
- *Cybersecurity and the Electric Grid | The State Role in Protecting Critical Infrastructure*, National Conference of State Legislatures Report (2020).

More information about the current state of DERs and interconnection standards can be found in *GMLC Survey of Distributed Energy Resource Interconnection and Interoperability Standards*, NREL and Lawrence Berkeley National Laboratory Report (2020).

Appendix F: Independent Review of the Proof-of-Concept Cyber100 Compass Cybersecurity Risk Tool

The independent review from Sandia National Laboratories begins on the following page.

Independent Review of the Proof-of-Concept Cyber100 Compass Cybersecurity Risk Tool

Gregory D. Wyss

February 2024

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

CONTENTS

LIST OF FIGURES

Independent Review of the Proof-of-Concept Cyber100 Compass Cybersecurity Risk Tool

By Gregory D. Wyss, Ph.D.
Distinguished Member of Technical Staff
Cyber Systems Security Research & Development Dept.
Sandia National Laboratories,[1] Albuquerque, New Mexico USA

1. OVERVIEW

The U.S. Department of Energy (DOE) Office of Cybersecurity, Energy Security, and Emergency Response (CESER), and Office of Electricity (OE) commissioned the National Renewable Energy Laboratory (NREL) to develop a method and tool to enable electric utilities to understand and manage the risk of cybersecurity events that can lead to physical effects like blackouts. This tool, called Cyber100 Compass, uses cybersecurity data elicited from cybersecurity experts, then incorporates that data into a tool designed to be usable by cybersecurity non-experts who understand the system itself. The tool estimates dollar-valued risks for a current or postulated future electric power digital control configuration, in order to enable utility risk planners to prioritize among proposed cybersecurity risk mitigation options. With the development of the Cyber100 Compass tool for quantification of future cyber-physical security risks, NREL has taken an initial bold step in the direction of enabling and indeed encouraging electric utilities to address the potential for cybersecurity incidents to produce detrimental physical effects related to electric power delivery.

As part of the Cyber100 Compass development process, DOE funded NREL to seek out an independent technical review of the risk methodology embodied in the tool. NREL requested this review from Sandia National Laboratories, and made available to Sandia a very late version of the project report, as well as NREL personnel to provide clarification and to respond to questions. This paper provides the result of the independent review activity.

The conceptual foundation of the method is based on work by Hubbard and Seiersen [Hubbard 2016 & 2023], which encourages quantification of cybersecurity risk in terms of potential future expected losses, expressed in dollars lost in one year, given known, expected and/or proposed conditions and practices for Operational Technology-related (OT) computing, networking and digital control infrastructure. The potential expected losses are identified as "Recommended Expectation of Loss" (REL) in order to underscore that these are merely recommendations suitable for planning purposes and not actual loss predictions. The method also asks utility executives to identify their risk tolerance in the form of likelihood of occurrence for particular dollar-denominated levels of loss. The tool uses Monte Carlo sampling to explore a large number of potential yearly scenarios and turn them into a statistical distribution of cybersecurity risk, presenting the losses for these scenarios to the user as a statistical exceedance graph for comparison to the identified risk

[1] This document has been approved for Unlimited Release as SAND2024-03315R. This work was performed for National Renewable Energy Laboratory with support for the work provided by the U.S. Department of Energy Office of Electricity (OE) under Agreement ID 39826 and Office of Cybersecurity, Energy Security, and Emergency Response (CESER) under Agreement ID 39915.

tolerance. Finally, analysts are able to perform "what if?" assessments to explore how changes in the OT infrastructure security conditions and practices will affect the cybersecurity risk distribution, enabling analysts, OT engineers, and C-Suite managers to have a productive dialogue about cost-effective steps toward managing cybersecurity risk.

Cyber100 Compass was developed to operate within the following scope and constraints, which are paraphrased here from Section 1.1 of the document for brevity:

1. Questions presented to the user must be answerable by people who are familiar with the OT infrastructure and operations but are likely not cybersecurity experts.
2. Given the information to be collected from the user above, the under-the-hood cybersecurity data required to turn that information into REL and other relevant metrics must be obtainable, either from data or from well-structured elicitation of subject matter experts (SMEs).
3. The tool should produce quantitative rather than qualitative risk estimates, preferably REL.
4. The framework should be extensible – it should be easy to incorporate additional OT cybersecurity conditions and practices.
5. The framework should address OT attacks that produce physical effects, not merely information effects.

2. NOTEWORTHY ELEMENTS OF THE METHOD AND REPORT

Target Audience of the Tool. It is particularly noteworthy that the target audience for the tool is for people who are not cybersecurity experts. That is, the tool incorporates the expertise of cybersecurity experts in the "back-end data", obtained via elicitation, with regard to the likelihood of particular cybersecurity events and their expected physical effects for a "typical" US electric utility. SMEs were instructed to think of a "typical" utility as one otherwise operating under normal conditions and in compliance with applicable federal and state regulations. The SMEs were also elicited regarding the degree to which the presence or absence of particular OT infrastructure security conditions and practices might cause the likelihood of event occurrence to increase or decrease, and to estimate a multiplicative factor by which a particular condition might change the default likelihoods (e.g., decrease its value to 80% of the baseline). The user benefits from this SME expertise by selecting which conditions and practices are present within their OT infrastructure without having to fully understand their significance to cybersecurity. The tool would then calculate the REL risk distribution. This approach would help cybersecurity novices to understand how particular OT security conditions and practices affect risk. The tool provides a starting point for risk mitigation and system improvement discussions with systems planners and executive-level decision makers. This approach is laudable.

References. The references cited in the report and used in the methodology development are reasonable and appropriate. The addition of a larger bibliography assures the reader that a larger context has been considered during the development of this methodology. The use of rubrics from NIST SP800-30 to understand the adversarial context of OT systems is noteworthy. For readers who are interested in gaining a better understanding of the challenges involved in adversarial risk assessment, a list of suggestions for additional reading is provided at the end of this Appendix.

Documentation of Condition Inputs. Documentation of the elicitation questions by which the SMEs were asked to provide data for the Condition inputs is welcomed. This documentation provides context for the breadth of condition inputs considered, the background and context provided to the reviewers, and the types of information that the software will expect users to provide. The breadth of coverage for the set of condition inputs is reasonable for an initial methods development effort, and the elicitation questions themselves are generally well-written.

Definition of Variables. The report provides solid definitions for the computational and elicited variables that explicitly represent conditional probabilities, event avoidance values (costs), risk tolerance, and REL, so that persons who either provide or consume these pieces of information should be able to come to a clear understanding of the actual meaning of each variable. The glossary provides a good summary and reminder of the meaning of terms in the context of this report.

Acknowledgment of Limitations. The document notes multiple times that this the method and tool should be viewed as "proof of concept" and not ready for deployment. The report itself does a good job of documenting what amount to internal critical review comments and observations in Chapters 7 (Known Issues) and 8 (Possible Follow-on Research), Appendices A (Next Phases of Work) and B (Subject Matter Engagement), as well as numerous comments throughout the document. These observations reflect good understanding of many of the warts in this initial method development activity.

Use of Sparse SME Data. The team made good use of the admittedly limited SME data they were able to elicit, but an obvious area for improvement is in the breadth SME engagement. Because of limited project time and especially limited availability from already-busy SMEs, the team was understandably able to get only four sets of data from which to base their back-end data. This means that the data is sparse, and in some cases, the distribution of data values is relatively broad. Use of the statistical mean among these data points means that all SMEs were considered to have equal weight, and all SME opinions affected the calculations equally. This is appropriate for such sparse data when there is no obvious reason to question the qualifications of any of the responding SMEs, although propagating the uncertainties identified in the elicitation would be more desirable (see related opportunity for improvement below).

3. OPPORTUNITIES FOR IMPROVEMENT

Editorial Revisions. During the review process several minor suggestions were made for editorial revisions to the text of the report, and were graciously accepted by the authors.

Use of Bernoulli Distribution. In Section 3.2.2 the method's developers model the likelihood of occurrence for each potential outcome (consequence type) using a Bernoulli distribution. For each simulation round (defined as a year), the Monte Carlo analysis method draws a random sample for each Bernoulli distribution to determine whether that specific outcome (e.g., equipment harm) occurs during that simulated year due to a cyber attack. The nature of the Bernoulli distribution means that the answer to this question is either Yes or No, i.e., either 0 or 1 events of this type occur during the simulated year. First, the Bernoulli distribution assumes that all events are independent, yet a first successful attack may well make subsequent attacks more likely (if same attacker can exploit the same vulnerabilities again) or less likely (if mitigation measures are taken). In addition, use of the Bernoulli distribution eliminates the very real possibility that more than one such event occurs during the year, especially for consequence-causing events which have a higher likelihood p_o. While it is recognized that the SME elicitation for p_o asked the SMEs to consider the likelihood that such an event would occur during a year (which would imply a Bernoulli distribution), this formulation may underrepresent risk by neglecting the possibility of multiple occurrences. Use of a Poisson distribution, with adjustment to the elicitation process to elicit an occurrence frequency rather than a probability, could overcome this limitation as it would produce a random variable for the number of times the event occurs during a year.

"Unsure" as a User Input Option. The report clearly describes the developers' intent for including an option in the user input to state that they are "unsure" about whether a particular condition exists in their OT cyber infrastructure. A condition, as used in Cyber100 Compass, is a term for "any constraints, resources, requirements, controls, or other factors that modify the cybersecurity risks of a system's energy transformation plans." As stated in Section 3.3, "users of Cyber100 Compass are not expected to have perfect knowledge regarding their energy systems… If a question asks for one and only one answer, Cyber100 Compass gives users the opportunity to select 'unsure'; however, users are encouraged to use this option as infrequently as possible." The text goes on to encourage analysts to select more decisive answers for the sake of the accuracy of the Cyber100 Compass results. The text also indicates that the specific condition for which "unsure" was selected will be neglected in the risk computation.

Recall that conditions *modify* the cybersecurity risk, either reducing or increasing that risk. Thus, for a utility where conditions exist that are unknown but detrimental to cybersecurity risk, those "unsure" entries will lead to a Cyber100 Compass risk results that underestimate the system's real risk, possibly leading to a false sense of security for the system owner and possibly errant risk management decisions. While the converse is also true (i.e., unrecognized conditions may exist that reduce cybersecurity risk can lead to falsely high risk results), this reviewer believes that the more dangerous situation likely leads to risk underprediction. As an example, consider:

- A hypothetical utility has low OT cybersecurity, possibly because they are early in their OT cybersecurity journey, and
- Their OT security personnel are still working on developing their expertise, and
- These personnel answer "unsure" as to their system's condition with regard to multiple questions, possibly because their training has not yet touched on the conditions described in the question, or they are otherwise not yet equipped to properly answer the question.

This could lead to significant underprediction of risk. Conversely, a utility that already has high-quality cybersecurity is likely to have that *because* they have well-trained cybersecurity personnel who understand both the security landscape and deeper details of their system's security conditions. Thus, the "unsure" option appears to present a higher likelihood for security risk to be underpredicted than overpredicted.

To ensure that analysts recognize the fact that the displayed risk results are uncertain when the "unsure" option has been selected, it is recommended that the tool's results screen show an explicit flag or warning whenever the user input data contains "unsure" values. Beyond that, it may be useful to provide the analyst with a computed upper bound for their security risk based on a presumed unfavorable resolution of the "unsure" values. However, computing "conservative" upper bound risk values often produces unrealistically or even laughably high risk estimates, so an opportunity for future investigation is to look for a method for computing a meaningful upper bound to be displayed for users when "unsure" data is present.

Calculation Method for Baseline Probabilities. Section 4.1 describes the method used for computing baseline probabilities. A baseline probability represents the likelihood that a specific adverse outcome will occur during a year within the 5-year planning horizon generally considered by the Cyber100 Compass tool.[2] Although not elicited as a frequency, it is closely related to the expected frequency of *successful* attack for each adverse outcome (successful meaning an attack that produces a physical manifestation for that specific adverse outcome). The methodology for calculating baseline probabilities has flaws in 1) implicit assumptions regarding the independence of variables and 2) implementation.

Cyber100 Compass decomposes each baseline probability into three component dimensions, expressed as multiplicative factors, in order to motivate thorough consideration by the SME being elicited:

- How would you rate the capabilities (skills and resources) of adversaries who seek to successfully execute this specific attack?
- How would you rate the likelihood that adversaries are motivated to target a particular electric grid system or component and execute this specific attack?
- How would you rate the severity of the vulnerabilities within a utility that an adversary might exploit to successfully execute this specific attack?

The definitions for the factors themselves as well as the rubrics used for the definition of the qualitative ordinal values "Very High," "High," and so forth are taken directly from NIST SP800-30, and from this perspective are reasonable and well-defined. The method elicits an ordinal value for each of the three factors, and uses the ordinal value to assign a numerical value over the [0-1] range. The factors are elicited separately for each adverse outcome. The numerical values are multiplied to obtain a candidate baseline probability for that outcome. The tool provides the analyst with initial quantitative dollar-valued estimates for annualized losses that the elicited values imply, and the SME is then encouraged to adjust their ordinal values to achieve a state in which the annualized expected losses and outage likelihoods are consistent with the SME's expert estimates.

While the thought process for this decomposition is sound as a method to motivate the thoughts of SMEs,[3] it is clear that these three factors are all interdependent with each other. However, the elicitation process treats them as independent, treats each [0-1] scale as a probability, and multiplies these values to obtain the baseline probability, implying a probabilistic logical AND operation among these dimensions. The use of a multiplication operation when combining these factors is only valid if the probabilities are either independent or conditional – a condition that is clearly not present in the definitions of the three factors. Furthermore, the definitions for two of the three dimensions do not indicate that a likelihood or probability is being elicited at all, and further digging into NIST SP800-30 shows that none of the qualitative ordinal definitions describe a likelihood or

[2] Given the speed with which OT cyber technology changes, and the speed with which new vulnerabilities in OT are discovered and exploited, using a 5-year planning horizon for REL may lead to inaccurate risk projections.

[3] Hubbard [Hubbard 2023] stresses the importance that elicited experts be "calibrated" prior to elicitation in order to ensure that potential quantitative biases are reduced and uncertainties properly considered. The Cyber100 Compass report does not indicate that calibration of experts was performed. This is a further opportunity for improvement.

probability for any of the three dimensions. The SME is asked to compensate for this by adjusting their qualitative answers until the annualized expected losses are consistent with their judgment. A method of this sort conflicts with good expert judgment elicitation methods, and the mathematics used (multiplying values from non-probabilistic factors to obtain a probability) are not consistent with probabilistic mathematics. A future independent review of this method (e.g., by the National Academies) will likely have strong objections to the current method. Methodology directions that may help resolve this are discussed in the section "An Area of Academic Controversy" found later in this review.

Consideration for Uncertainty of Probabilities. In the Bayesian view of probability that is found throughout the Cyber100 Compass methodology, probability can represent a person's degree of belief regarding the potential for future events to occur. However, the uncertainties in our understanding of these probabilities leads to uncertainties in the risk results, both with regard to the magnitude of the computed risk and the ranking of specific scenarios' risks. That is, the ordering of scenarios within a rank-ordered list according to priority for risk-informed mitigation is itself uncertain, and this fact is disguised when point estimate values for probabilities are used without their associated uncertainties in risk computations. Cyber100 Compass uses only point estimates for probabilities throughout, so decision makers are given a false sense of confidence regarding the rank ordering of the importance of individual scenarios for risk mitigation. The tool also does not present an analyst with confidence intervals for the REL values it computes, and in modern risk studies, analysts consider the confidence intervals for the results to be almost as important as the results themselves. This issue becomes more and more severe as the scenarios considered in the risk assessment grow more numerous and specific, so the high-level approach to risk assessment used in Cyber100 Compass may be somewhat less susceptible than many others to the rank reversal phenomenon. However, including such uncertainties in the risk computations is an important opportunity for improvement.

When considering uncertainty, methodology developers should note that there is a negative side effect of including uncertainty in these calculations. When the rank ordering of scenarios for mitigation is itself uncertain, it makes decision makers' task of selecting the most cost-effective mitigation measures more difficult and uncertain. It is an unfortunate feature of security risk assessment writ large that methods developers are faced with the choice of presenting results that are uncomfortably uncertain or results that likely contain rank reversals because the analyst's beliefs about the likelihood of specific scenarios does not match actual but poorly known adversary decision making preferences. Practically speaking, one method to compensate for these uncertainties is to lean toward resilience measures for mitigation solutions – measures that one would expect to be broadly applicable across a wide range of scenarios, including natural, malevolent, and human error. Mitigations that are closely tied to a very specific scenario have the potential to merely push the adversary to pursue a different and less-protected scenario, a phenomenon known as "threat shifting".

In situations where humans are deliberately plotting malevolent behavior, the dominant probabilistic uncertainty is generally the likelihood that the adversary decides to make any attack at all – i.e., the problem of quantifying deterrence. This is discussed at greater length in the next section.

4. AN AREA OF ACADEMIC CONTROVERSY

The two most significant opportunities for improvement described above point toward a larger area of ongoing academic controversy – one that has far-reaching practical impacts. Stated briefly, the question is "How should the likelihood of attack be considered in security risk assessment?" The root of the problem lies in the fact that initiating an attack is a deliberate malevolent decision that is the culmination of a planning process that involves the adversary's own cost-benefit-uncertainty-risk assessment according to their own value systems, which are often poorly known to us and can change rapidly as driven by unfolding world events. The implications of this issue reach throughout the security risk analysis and management discipline and have resulted in significant academic disagreement. This section first describes two opposing positions for treatment of this issue, then discusses how the answer to this issue may be different for some adversary types vs. others. The section closes with a discussion of how Cyber100 Compass methods might evolve to restructure the elicitation and computation of p_o and $p_{L/M/H}$ as well as implications for uncertainty analysis. The Suggestions for Further Reading at the end of this review provide opportunities for readers to dig deeper, primarily on these topics.

4.1. The Debate

On one side of this debate are Hubbard and Seiersen (authors of the book on which Cyber100 Compass draws heavily), along with former president of the Society for Risk Analysis (SRA) Barry Ezell and many others. This side of the debate holds that the likelihood of attack can and should be treated as a Bayesian probability or frequency because the elicited value represents our best current understanding of that likelihood. They acknowledge that the uncertainty in these elicited values is large – even several orders of magnitude for terrorists attacks that require very high adversary capabilities, but that placing the risk results within the context of these large uncertainties is the proper way to provide decision makers with the full picture of what is actually known about the risk landscape. They acknowledge that the high uncertainty can be challenging for decision makers to deal with, and emphasize that the values should be updated or re-elicited as additional information relevant to the attack likelihood emerges – just as one would do for any elicited Bayesian likelihood. The Cyber100 Compass tool is fully aligned with this viewpoint.

The other side of this debate is championed most notably by Tony Cox, the former editor of the SRA journal *Risk Analysis*, although he is joined by many others. They do not dispute the claims of the Ezell-Hubbard camp, but rather point out how the planning and decision making process of malevolent humans introduces nonprobabilistic elements that are important to security risk *management*. In particular, they argue that the decision to attack is not random, and is especially not an independent random variable, as compared to the independence of the initiating event in a safety risk analysis. Indeed, a likelihood of attack is possibly the most dependent variable in the risk equation, as an adversary considers a scenario from the perspectives of whether the expected outcome of the attack will be satisfactory, whether the resources required to perform the attack are available, whether the expenditure of these resources is worthwhile to attain the expected outcome, whether the likelihood of attack success is acceptably high, whether there is an acceptable degree of personal or organizational risk engendered by attempting to carry out the attack, and especially whether there are any "better" options available to achieve the desired outcomes (as such, this clearly cannot be an independent variable!). Thus, when an attack likelihood is directly elicited as a Bayesian variable, the SME must have all of these factors in mind if they are to produce an accurate likelihood and uncertainty. The listed dependencies lead to adversary behaviors of *adaptation* (an adversary adapts, i.e., acquires all of the additional capabilities and information required to achieve a successful attack), *threat shifting* (an adversary chooses a different attack path, a different target, etc.), and *deterrence* (an adversary decides, for any reason, not to attempt any attack). Risk assessments that cannot include these behaviors, they argue, cannot accurately model security risk because these nonprobabilistic behaviors are fundamental to the security risk landscape. Indeed, in physical security, an important goal of many upgrades is to influence an adversary's decision making process in favor of deterrence or threat shifting. And at best, an elicited attack likelihood requires the SME to account for such behaviors in the elicitation process rather than making them an explicit element of the risk analysis. Cox has argued these behaviors dictate that the likelihood of attack should be an output of a risk assessment, not an input to it.[4] [Cox 2009]

Eliciting a likelihood or probability of attack requires consideration of all of the elements that affect the likelihood, some of which often go unrecognized. Figure 1 was developed by this reviewer [Wyss 2022] to show an extended but possibly still incomplete list of the dependencies that affect an attack likelihood, especially as it relates to a high-level adversary executing a highly advanced attack scenario. It can be challenging for experts to consider an appropriate set of these when formulating their likelihood estimates, especially for attacks by higher-level adversaries and for insider attack scenarios, where specifics of motivation and intent weigh so heavily in that likelihood.

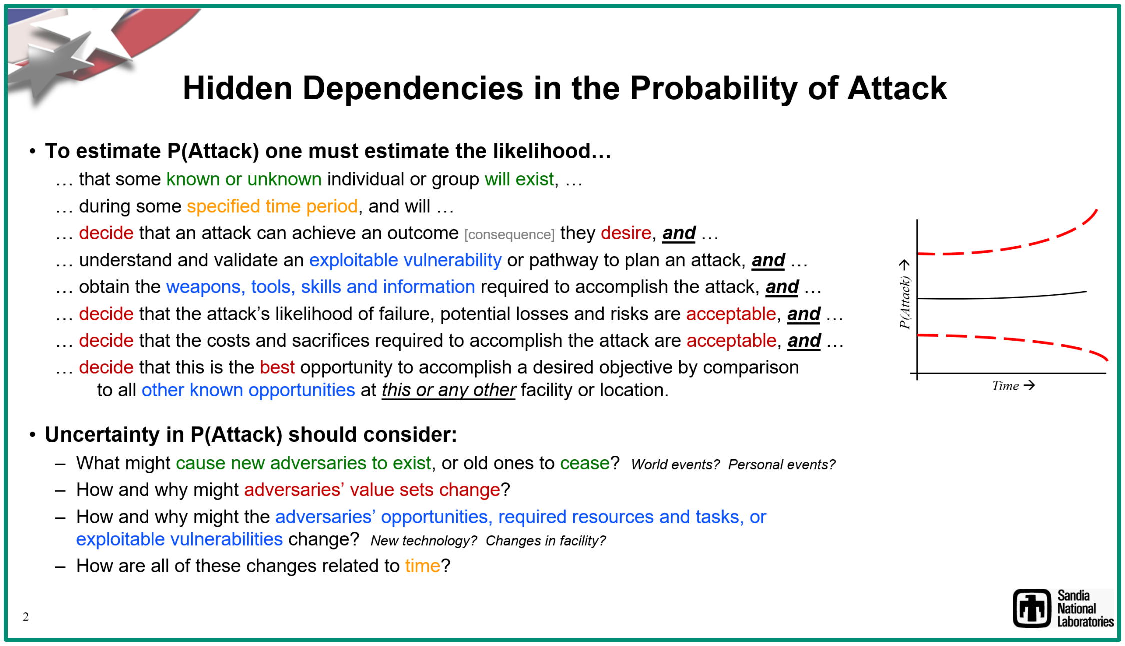

Hidden Dependencies in the Probability of Attack

- **To estimate P(Attack) one must estimate the likelihood…**

  … that some known or unknown individual or group will exist, …
  … during some specified time period, and will …
  … decide that an attack can achieve an outcome [consequence] they desire, ***and*** …
  … understand and validate an exploitable vulnerability or pathway to plan an attack, ***and*** …
  … obtain the weapons, tools, skills and information required to accomplish the attack, ***and*** …
  … decide that the attack's likelihood of failure, potential losses and risks are acceptable, ***and*** …
  … decide that the costs and sacrifices required to accomplish the attack are acceptable, ***and*** …
  … decide that this is the best opportunity to accomplish a desired objective by comparison to all other known opportunities at *this or any other* facility or location.

- **Uncertainty in P(Attack) should consider:**
  - What might cause new adversaries to exist, or old ones to cease? *World events? Personal events?*
  - How and why might adversaries' value sets change?
  - How and why might the adversaries' opportunities, required resources and tasks, or exploitable vulnerabilities change? *New technology? Changes in facility?*
  - How are all of these changes related to time?

Figure 1. The likelihood of attack is arguably the most dependent variable in a security risk assessment, and has many hidden dependencies. Dependencies and uncertainties shown in green relate to adversary existence; those shown in red relate to adversaries' value sets; those shown in blue relate to adversaries' capabilities and opportunities; and those noted in orange relate to time. The graph notionally denotes how uncertainty increases with the passage of time because any of these can change for possibly unpredictable reasons as the time horizon becomes longer. [Wyss 2022]

[4] The intensity of this debate should not be underestimated. For example, the literature includes a significant exchange between Ezell et. al. and Brown and Cox on this topic in the journal *Risk Analysis*, which is listed as a group in the "Cited References" at the end of this review.

4.2. The Effect of Adversary Characteristics

Some, including this reviewer, believe that the specific characteristics of the adversary and attack scenario may dictate which of these viewpoints is most appropriate to a specific situation. For this discussion it is useful to view the spectrum of adversaries and attacks through the lens used by the Defense Science Board (DSB), who described adversaries in terms of six "Tiers" as shown in Figure 2 [DSB 2013]. They describe these Tiers briefly as follows:

- Tiers I and II attackers primarily *exploit known* vulnerabilities
- Tiers III and IV attackers are better funded and have a level of expertise and sophistication sufficient to *discover new* vulnerabilities in systems and to exploit them
- Tiers V and VI attackers can invest large amounts of money (billions) and time (years) to *actually create* vulnerabilities in systems, including systems that are otherwise strongly protected.

Higher-tier competitors will use all capabilities available to them to attack a system but will usually try lower-tier exploits first before exposing their most advanced capabilities.

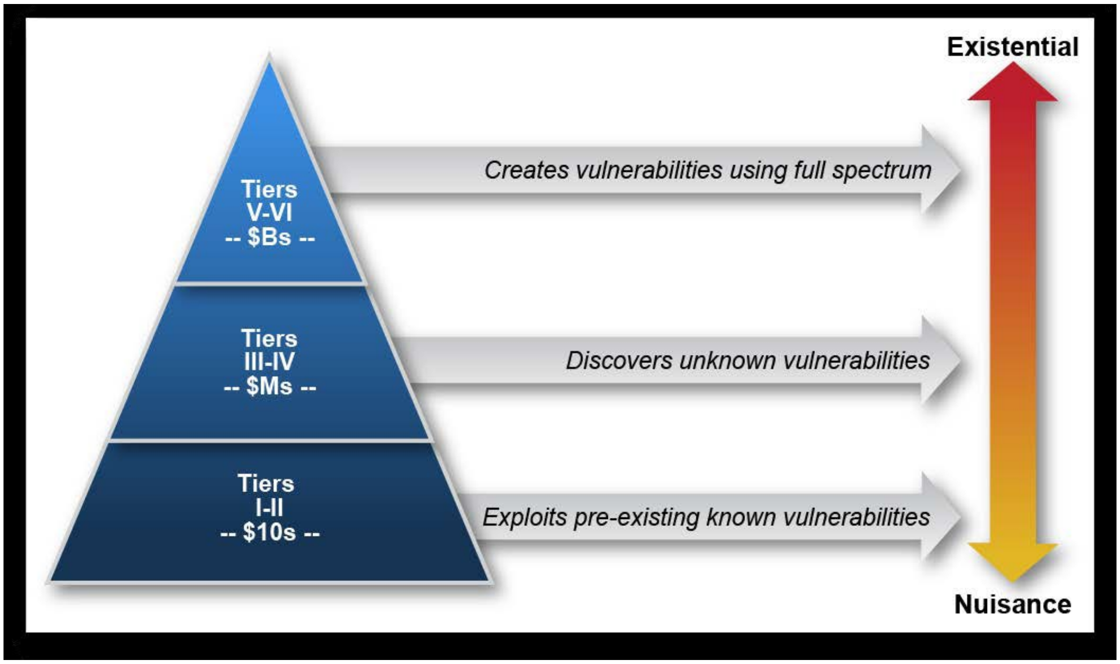
Existential
Tiers V-VI -- $Bs --
Creates vulnerabilities using full spectrum
Tiers III-IV -- $Ms --
Discovers unknown vulnerabilities
Tiers I-II -- $10s --
Exploits pre-existing known vulnerabilities
Nuisance

Figure 2. Cyber Threat Taxonomy as described by the Defense Science Board [DSB 2013]

A low-level (Tier I) adversary uses tools created by others to exploit already-identified vulnerabilities, and their opportunities are limited to systems using equipment that contain vulnerabilities that can be exploited by readily obtainable tools. Access to digital systems is obtained by common network exploits or burglary techniques. The opportunities for attack and the obtainable consequences are limited by the available targets and tools. Such adversaries do not generally invest large amounts of money or time to prepare and execute attacks, and they typically choose attacks for which there is little expectation of attacker risk (e.g., attribution and punishment for the attack). These "attacks of opportunity" are the most common attacks, or attacks that can be extrapolated from common attacks. Evidence is available to inform estimates of scenario likelihood based on occurrence rates for similar attacks or attacks on comparable targets, so these types of adversaries and scenarios are the most compatible with the Ezell/Hubbard position, and by extension, Cyber100 Compass.

A high-level (Tier VI) adversary often represents a nation state or other high-capability research-based entity, and is able to *create* vulnerabilities in systems, discover novel vulnerabilities in existing equipment, use all manner of existing attack tools, and develop new attack tools. The highest-level adversaries have means to gain clandestine physical and electronic access to most systems, bypassing or neutralizing state-of-the-art security systems. This adversary type typically selects targets on the basis of the desired outcome (type, magnitude, location, etc.), and while they try to minimize their attacker risk, they do not shy away from such risk when it is deemed necessary to achieve their desired outcome. Such an adversary often spends years and uses a large staff to develop and execute these attacks. The opportunities for attack are limited only by known science, bolstered by research, and can be targeted to attack almost any entity. These "targeted attacks" by the highest-level adversaries are much more rare, and can be unprecedented until they are discovered. Examples of such attacks on OT include Stuxnet and the 2015 cyber attack on Ukraine's electric power grid. The most important contributor to the likelihood of such attacks, where they are possible, is whether a high-level adversary develops a desire to inflict a particular consequence on a specific target or population of targets. Much of the evidence for future occurrence of such attacks necessarily comes from the intelligence community, and such evidence is often highly uncertain in target identification, likelihood, and attack objectives. Thus, these types of adversaries are more compatible with the position of Cox et al, and as such cannot be represented very well in Cyber100 Compass.

The DSB describes 6 tiers of adversaries, bounded by the above examples. As the level of adversary increases within this range, the degree of compatibility with the Ezell/Hubbard perspective decreases, and compatibility with the Cox perspective increases. Uncertainty in attack likelihood also increases significantly with adversary level. For these reasons, the REL results for Cyber100 Compass should be viewed as less uncertain for attacks of opportunity and more uncertain or even speculative for targeted attacks. The authors should consider noting this issue as a limitation to the methodology in the report.

Attacks by insiders using their knowledge, access and authority, combined with relatively low-level attacker tools, can achieve attack outcomes generally associated with much higher-level outsider attacks. Therefore, if one estimates attack likelihoods, the elicitees must also include all forms of insiders in their likelihood estimates. This is a significant challenge because of the high variability in motivation and intent of human insiders. In addition, high-tier actors are able to operate deep within OT systems, gaining attack capabilities generally limited to insider system operators with high privileges. In many cases, well-placed electronic subversions can cause outcomes similar to human insiders, leading to the humorously-framed observation that systems need to defend against both "carbon-based insiders" (humans) and "silicon-based insiders" (electronic subversions).

5. POSSIBLE PATHS FORWARD

The most important need for improvement in Cyber100 Compass is in elicitation of p_o, as a future independent review of this method (e.g., by the National Academies) will likely have strong objections to the current method. The decomposition of p_o into three factors is useful for motivating the thoughts of experts, but the method by which each qualitative ordinal descriptor is assigned a quantitative probability is not identified or validated in the document. Furthermore, these factors are interdependent, so multiplying them as probabilities is not appropriate. But just as important, the three factors elicited for p_o are at a level of detail that is finer than that for which robust data is available. Hubbard [Hubbard 2023] indicates that elicitation should be done at a level consistent with data availability, and for attacks of opportunity, data is most available at the level of p_o itself. Information specific to cyber attacks leading to physical effects for electric power grids may be sparse, but other sources of information are available that can inform the estimate in a Bayesian sense, for example:

- Electric power attacks where an adversary achieved access sufficient to cause a physical outcome had they decided to do so ("precursor events"),
- Attacks on other critical infrastructures' OT systems that led to or could have led to physical outcomes, and
- Attacks on OT infrastructures in other domains such as manufacturing.

Such estimates become less confident and contain greater uncertainty for targeted attacks by higher-level adversaries and for insider attack scenarios because of the dearth of available data and because of the high variability in motivation and intent of human insiders. Elicitation of *uncertainty* in the p_o values should also be an integral part of the elicitation process.

A more ambitious path forward would be to adapt the Cyber100 Compass risk model so that its primary analysis is for conditional risk. A conditional risk analysis is statistically conditional upon a successful attack having occurred, sometimes colloquially stated as "we set the likelihood of a successful attack to 1.0" to examine the distribution of outcomes, pathways and consequences. The most straightforward mathematical formulation for a conditional risk analysis, given the current Cyber100 Compass risk formulation, is to separate the current variable p_o into two parts, so $p_o = p_A \cdot p_{En}$. Here p_A is the likelihood that an attack occurs *and* that this attack leads to at least one undesirable physical effect, and p_{En} represents the conditional likelihood that the nth undesired physical effect occurs given this successful attack. The Cyber100 Compass report describes five undesired physical effects in Table 1, listed as Power Outage, Harm to Equipment, Harm to Employees, Harm to Community, and Loss of Productivity or Efficiency, which might be thought of as n=1 through n=5 for this example.

The benefits and drawbacks of the conditional risk approach are as follows. Neglecting p_A (or colloquially, setting p_A to 1.0), which leads to a computation of conditional risk, has the effect of avoiding the "what is the likelihood of attack even by the highest-level attacker?" question. Thus it sidesteps the academic controversy described previously. In addition, if one wants to perform an unconditional risk assessment, the mathematics are clear, and the characteristics of the quantity being elicited for p_A are clearly and explicitly defined as embodying *all* the various types of uncertainties related to likelihood of attack. Furthermore, the other terms in the equation may be less uncertain with respect to adversary motivation and intent since the dominant uncertainties are contained in p_A. However, mathematics rarely lets us off the hook with a "free lunch," and the conditional risk approach appears to bury some of the complexities caused by the multiplicity of adversary types in the elicitation of the Impact Scale and Impact Level Probabilities, as shown in notionally in Figure 2 and specifically in Figure 5 of the report. Other decompositions of the problem are possible, and it would be useful for the methods development team to consider a range of possibilities in this regard.

The topic of uncertainty looms large over most security risk assessments. The Cyber100 Compass tool explicitly samples over a range of possible consequence magnitudes for each Monte Carlo trial, and this is methodologically good. The team would do well to validate the distributions that result from the convolution of the selection among the high/medium/low consequence level and the Monte Carlo sampling done within each consequence level. Also, the tool does not consider the uncertainties associated with any of the probabilities included in the model, and these are almost certainly large. Incorporating uncertainty in these probabilities into the tool would be a major undertaking, but providing REL results without uncertainties leads to a significant possibility for rank reversals within the prioritized list of risks and poorly-informed risk decision making. At a minimum, the team should research ways to compute and display for the analyst a confidence interval for the results to indicate that "risks with computed REL values that differ by less than *x* should be considered equivalent," *i.e.*, statistically indistinguishable. An example of such a statement in another domain is the "margin of error" value listed by analysts for public opinion polls.

At the risk of stating the obvious, the Cyber100 Compass risk assessment method deals only with cybersecurity risks. Cyber attacks have unique characteristics, are of growing sophistication and frequency, and can cause increasingly severe physical effects as digital control systems penetrate more pervasively into grid operations. However, risk analysts should be aware that cyber attacks are only one facet of the multidimensional security risk landscape for the electric grid. If a determined adversary can achieve their objectives more easily or with greater duration by other means, such as physical disablement or destruction of critical power-handling equipment, they may well select that attack pathway even though a viable cyber attack pathway exists. Utility risk managers should be reminded that security risk management must be balanced across all attack domains lest the result be a system with the proverbial "firmly locked doors but wide-open windows" problem. It was stated earlier that resilience measures (measures that are broadly applicable across a wide range of scenarios, including natural, malevolent, and human error) are a practical method to compensate for the large uncertainties associated with cybersecurity risks. But diversity in resilient mitigation measures, to include elements of both cyber and physical mitigation pathways, can provide grid operators with the maximum set of options to respond to all types of physical outcomes, potentially even reducing consequence duration and severity for attacks by high-level adversaries using unprecedented attack vectors.

This final point is broader than Cyber100 Compass. Given the extraordinary difficulty of risk quantification and risk management decision making for high-tier adversaries, risk analysts may wish to consider other complementary risk management approaches for these adversaries. One approach that has seen some success in the physical security arena has been a nonquantitative method where risk analysts prioritize attack scenarios for mitigation not on the basis of a highly uncertain quantitative risk but rather on the basis of the relative difficulty an adversary would experience while planning and executing the attack [Wyss 2013]. Elements of "scenario difficulty" can include:

- the need to acquire or develop advanced tools or weapons,
- the need to cultivate advanced or rare skills,
- the need to obtain and exploit closely-held information, and to verify its authenticity,
- the potential for unrecoverable errors or task failures to occur,
- the need to act within a possibly narrow time window, and

- exposure of the attacker or their support group to possibly unacceptable personal risk.[5]

A high-tier adversary's attack planning process looks across all of these elements, and possibly others, to identify the possible stumbling blocks within the attack pathway. Among those attack pathways that meet their criteria for an "acceptable outcome", the most attractive attacks are those that have the fewest and least significant disadvantages with regard to these stumbling blocks, and *not* those with the most "advantages," because even a single insurmountable stumbling block renders all of the supposed advantages irrelevant. Attack pathways can be prioritized by defenders for mitigation according to their difficulty and consequences, under the realization that pathways that are "easier" and lead to higher consequences represent higher security risks and should thus be a higher priority for mitigation when compared with those that are inferior in these metrics (*i.e.*, prioritize non-dominated scenarios for mitigation as they are expected to be more attractive to adversaries). This method has been applied most successfully to the physical security domain, but the thought process is useful when prioritizing among cybersecurity pathways with regard to risk mitigation decisions, particularly with regard to high-tier adversaries.

6. SUMMARY

This paper documents the results of an independent technical review of the Cyber100 Compass cybersecurity risk management method and tool developed by the National Renewable Energy Laboratory. The method was developed for assessment of cybersecurity risks as they relate to potential cyber-induced undesired physical effects on electric utilities and the electric power grid, and how those risks may change over time as renewables become an increasing presence on the grid and as new cybersecurity features are added to its digital control systems. The method relies heavily on the elicitation of data from experts to populate its back-end data in order to enable non-expert analysts to get meaningful results from the tool. Noteworthy elements of the method and report identified during the review include:

- The target audience for the tool is people who are not cybersecurity experts, which helps make the tool more broadly useful across the electric power domain,
- The report provides good definitions for the computational and elicited variables that used in the method for persons who either provide or consume this information should be able to clearly understand the meaning of each variable,
- The documentation provided to experts from whom data is elicited was found to be well-written and useful for the elicitation process,
- The breadth of coverage for the cybersecurity conditions presented to the experts for their consideration was found to be reasonable for an initial methods development effort,
- While the back-end data that the team was able to collect from experts was sparse, the team made good and appropriate use of the data they could obtain,
- The document provides a thorough introspection of the method, acknowledging its limitations throughout, and

[5] "Unacceptable risk" for a cybersecurity attacker could be as simple as discovery and attribution at any time during the attack, or at least prior to the realization of their desired physical outcomes in the case of cyber-physical systems like the electric grid.

- The references cited in the report and used in the methodology development are reasonable and appropriate.

The review also noted the following opportunities for improvement:

- Use of distributions beyond the Bernoulli distribution for likelihood of physical outcomes,
- Revision of the mathematical treatment of the "unsure" option in the user input to reduce the potential for underestimation of risk,
- Revision of the method for eliciting and computing baseline probabilities, and
- Inclusion of uncertainties for probabilities within the methodology.

The review also provided context for an important ongoing academic controversy regarding likelihood of attacks. This controversy influences all security risk assessment methods. Suggestions for possible paths forward for the Cyber100 Compass method in light of this controversy are provided. It is acknowledged that the path forward for security risk computation writ large is uncertain in light of the noted controversy, and other methods for quantifying and managing cybersecurity risk are likely to be developed as this controversy continues to be debated and, one would hope, eventually resolved.

ABOUT THE AUTHOR

Gregory Wyss, PhD, Distinguished Member of Technical Staff
Cyber Systems Security Research and Development, Sandia National Laboratories

Dr. Wyss has worked in the areas of risk, reliability and vulnerability assessment at Sandia National Laboratories for over 35 years. During that time, he has performed risk assessment studies for nuclear reactors, space vehicle launches, nuclear assets and test facilities, telecommunications facilities, and a variety of other high-integrity and potentially high-consequence systems. He has also developed risk and vulnerability assessment and uncertainty analysis methodologies for a broad range of applications. He has expertise in the areas of security vulnerability analysis and risk analysis, safety risk assessment for space nuclear systems, system-theoretic process analysis, system-of-systems studies, and integrated risk uncertainty analysis using Latin hypercube sampling and other techniques. He has taught short courses on these subjects for more than 30 years.

Dr. Wyss has responsibilities that include development and implementation of cybersecurity risk management methods and requirements for military nuclear systems. His responsibilities also include development of enterprise security risk management methodologies for high-security nuclear facilities. His research interests include assessment of microelectronics supply chain attacks, insider threats, decision support, pre-attack planning, and synergistic effects between cyber and physical security systems.

Dr. Wyss holds a Bachelor of Science degree with Highest Honors in Nuclear Engineering from the University of Illinois at Urbana-Champaign (1983), and Master of Science and Doctor of Philosophy degrees from the University of Illinois at Urbana-Champaign (1985, 1987). He is a Fellow of the International Association for the Advancement of Space Safety.

REFERENCES AND SUGGESTIONS FOR FURTHER READING

Cited References

[Cox 2009] Cox, Jr., L.A., "Response: Game Theory and Risk Analysis," *Risk Analysis,* **29** (2009) No. 8, pp. 1062-1068.

[DSB 2013] Defense Science Board, US Department of Defense, *Task Force Report: Resilient Military Systems and the Advanced Cyber Threat*, US Defense Science Board Task Force on Resilient Military Systems and the Advanced Cyber Threat, January 2013.

[Hubbard 2023] Douglas W. Hubbard and Richard Seirsen, *How to Measure Anything in Cybersecurity Risk, 2nd Edition*, John Wiley and Sons, New Jersey, 2023.
Note: 1st edition, cited in the Cyber100 Compass report, was published in 2016.

[Wyss 2022] Slide taken from: Wyss, G.D., "Risk Informed Management of Enterprise Security (RIMES), or 'Why computing security risk based on a probability of attack is possible, but is likely not useful for Risk Management'", SAND2022-8311 C, Prepared at Sandia National Laboratories, Albuquerque, NM, and presented at PSAM 16, Honolulu, HI, June 2022.

[Wyss 2013] Wyss, G.D., Clem, J.F., Darby, J.L., Dunphy-Guzman, K., Hinton, J.P., and Mitchiner, K.W., "A Method for Risk-Informed Management of Enterprise Security (RIMES)", SAND2013-9218P, Sandia National Laboratories, Albuquerque, New Mexico, 2013.
Note: this is a reprint of the conference paper "Risk-Based Cost-Benefit Analysis for Security Assessment Problems" (SAND2010 5095C) and presented at the IEEE 44th Annual International Carnahan Conference on Security Technology, San Jose, CA, October 5-8, 2010.

The Exchange of Opinions Between Ezell et. al. and Brown and Cox, Listed Chronologically
Ezell, B.C., *et. al.*, "Probabilistic Risk Analysis and Terrorism Risk," *Risk Analysis*, **30** (2010) No. 4, pp. 575-589.

Brown, G., and Cox, Jr., L.A., "How Probabilistic Risk Assessment Can Mislead Terrorism Risk Analysts," *Risk Analysis*, **31** (2011) No. 2, pp. 196-204.

Ezell, B., and Collins, A., "Letter to the Editor," *Risk Analysis*, **31** (2011) No. 2, p. 192

Brown, G., and Cox, Jr., L.A., "Response: Making Terrorism Risk Analysis Less Harmful and More Useful: Another Try," *Risk Analysis*, **31** (2011) No. 2, pp. 193-195.

Other References

A broad discussion of the issues affecting risk management writ large
Douglas W. Hubbard, *The Failure of Risk Management: Why it's Broken and How to Fix it,* John Wiley and Sons, New Jersey, 2009.

A broad discussion of the issues affecting quantification of terrorism risk

Committee on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis, National Research Council, "Department of Homeland Security Bioterrorism Risk Assessment: A Call for Change." ISBN: 0-309-12029-2, National Academies Press, Washington, DC, 2008.

NIST references useful for cybersecurity risk management relevant to Operational Technology (OT) systems

Stouffer, K. , *et al*, *Guide to Operational Technology (OT) Security*, Special Publication (NIST SP 800-82 Rev. 3), National Institute of Standards and Technology, Gaithersburg, MD, [online], https://csrc.nist.gov/pubs/sp/800/82/r3/final (Accessed December 7, 2023)

Ross, R. (2018), *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy*, Special Publication (NIST SP 800-37 Rev. 2), National Institute of Standards and Technology, Gaithersburg, MD, [online], https://doi.org/10.6028/NIST.SP.800-37r2 (Accessed December 7, 2023)

Journal Papers, Listed Chronologically

Cox, Jr., L.A., "Does Concern-Driven Risk Management Provide a Viable Alternative to QRA?," *Risk Analysis*, **27** (2007) No.1, pp. 27-43.

Cox, Jr., L.A., "Some Limitations of Risk = Threat × Vulnerability × Consequence for Risk Analysis of Terrorist Attacks," *Risk Analysis*, **28** (2008) No.6, pp. 1749-1761

Insua, D.R., Rios, J., and Banks, D., "Adversarial Risk Analysis," *Journal of the American Statistical Association*, **104** (2009) No. 486, pp. 841-854.

Cox, Jr., L.A., "Some Limitations of Frequency as a Component of Risk: An Expository Note," *Risk Analysis*, **29** (2009) No. 2, pp. 171-175.

Cox, Jr., L.A., "Improving Risk-Based Decision Making for Terrorism Applications," *Risk Analysis*, **29** (2009) No. 3, pp. 336-341.

Cox, Jr., L.A., "Perspective: What's Wrong with Hazard-Ranking Systems? An Expository Note," *Risk Analysis*, **29** (2009) No. 7, pp. 940-948.

Merrick, J., and Parnell, G.S., "A Comparative Analysis of PRA and Intelligent Adversary Methods for Counterterrorism Risk Management" *Risk Analysis*, **31** (2011) No. 9, pp. 1488-1510.