OLIVER WYMAN

TAMING CYBER

QUANTIFYING CYBER RISK USING A STRUCTURED SCENARIO APPROACH

JANUARY 2018

AUTHORS

Ramy Farha, Partner

Evan Sekeris, Partner

Jerry Wu, Engagement Manager

MARSH GUY CARPENTER MERCER OLIVER WYMAN

MARSH & McLENNAN COMPANIES

At the heart of risk management is a gloomy truth: You can never achieve zero risk. In a world of limited resources there are always tradeoffs to be made: how much to invest here and how much there, how much risk to tolerate and how much to mitigate or insure against.

To answer those questions, risk quantification is necessary – to estimate how likely an outcome is to occur and more importantly, what will the cost be; translating complex real-world events into dollar figures that can enable rational decision making is critical to effective risk management.

Organizations understand this paradigm. Businesses, especially in financial services, are built on a foundation of assessing and comparing risk. But talk to a C-suite executive today, and you are likely to hear: "Cyber risk is one of our biggest concerns. We have experts who understand our systems and our data and who try to protect the organization." We think the most common misconception about Cyber risk and Cyber attacks is the perception that these attacks are purely technical – machines attacking machines. In practice, attackers rely heavily on understanding of people, policies, and how a company is organized – people attacking people. A fully hardened server is hopeless in the face of an employee who is tricked into opening a door to an intruder. Therefore, often times the C-suite concludes: "In terms of quantifying risk, we are in the dark. We do not know our true Cyber exposure. We cannot manage Cyber risk properly because we cannot measure the risk. We do not know how to best invest in risk mitigation."

Clearly, identifying and quantifying Cyber risk is different from quantifying "financial" risks (e.g., credit, market, etc.), and offers some unique challenges – especially the lack of data and the speed with which would-be attackers discover new vulnerabilities and devise new ways to exploit these vulnerabilities. To fully understand and quantify Cyber risk, one needs to understand technical and nontechnical avenues of attack.

In our work with clients over the past few years, we have seen the potential to substantially improve the process of quantifying Cyber risk, producing results that, while not perfect, can improve the ease and quality of executive decision making related to Cyber risk. Such quantification allows for better, more accurate answers to some of the key questions about Cyber risk that Boards and senior management teams need to answer: What is the overall level of Cyber risk exposure of the institution? Which areas of the institution drive the most material Cyber risks? What level of exposure to Cyber risk are we willing to accept? How much should we invest in Cyber risk mitigation – and where should the investments go? How much insurance should we purchase?

THE CHALLENGE OF CYBER

Our focus is on quantifying Cyber risk by assigning a dollar value to the costs to the organization of a successful attack. But we believe the new process of quantification we outline should provide far more than the answers to the question of how much an attack would cost. In our experience, the risk quantification exercise provides valuable insight into the role of data and digital processes in the institution, including any links to customer confidence, regulatory compliance, civil liability, and competitiveness. Cyber crime thrives in the gaps between disciplines. Therefore, the approach brings experts on both sides of the IT/Risk gulf to a sharper understanding of the connections between the respective disciplines, and contributes enormously to a culture of risk awareness across the organization. By doing so, the approach explicitly adds multiple dimensions to the quantification and allows for more effective risk management.

Many institutions today assess Cyber risk by soliciting the opinions of in-house IT experts on the topic. In our experience, these exercises are sub-optimal given the answers are more guesstimates rather than proper risk quantification exercises. As a result, the outcomes are relatively ineffective risk management tools. The weaknesses in these exercises should not be surprising, particularly in light of the following reasons.

Institutions lack historical data. Cyber risk is an emerging risk with limited useful historical data. And the situation is unlikely to change soon, because institutions are often unwilling to disclose the details of successful attacks, and especially the true cost of incurred losses.

Threat environment is rapidly changing. Attackers are constantly finding new ways to access the IT systems and infrastructure. What an institution knows about current vulnerabilities today is likely to become obsolete tomorrow. Without a structured process, institutions will find the task of keeping up with these changes extremely difficult.

Cyber attack outcomes are not always comparable. The impact and cost of various Cyber risk events such as a data breach or disruption are typically unique to the institution and highly dependent on the individual operational, IT, system, or data environment.

For these reasons, we believe that the best way to assess Cyber risk is through a structured scenario approach which combines quantitative and qualitative techniques focusing on the mechanisms that lead to losses to quantify the potential loss exposure. The approach is in contrast to a number of Value-at-Risk approaches used to quantify non-financial risks which typically rely exclusively on historical loss data. The process allows institutions to identify the most crucial assets (e.g., critical people, processes, data, and technology), greatest vulnerabilities, and likeliest avenues of attack, and provides a disciplined way to explore the full range of financial consequences from a potential event, providing further insights that can be leveraged for risk management purposes.

ADAPTING A POWERFUL (AND FAMILIAR) TOOL

Financial institutions, of course, are familiar with the idea of structured scenario analysis. The idea is key to the stress testing exercises mandated by US regulators since the 2008-09 financial crisis. While the scenarios used to analyze Cyber risk can be less detailed than the scenarios used for stress testing, the ultimate goal is the same – to tease out unexpected risks and render these risks more manageable to institutions by attaching dollar values to outcomes.

The Cyber risk quantification process using structured scenario analysis involves two broad phases: (i) identify Cyber risks and (ii) quantify Cyber risks. Risk identification is performed by first identifying high-value assets, and then identifying the risks associated with these assets. Risk quantification requires estimating both the frequency and severity of Cyber risk events. These estimates are convoluted to obtain a loss estimation for the events, and can be aggregated to obtain an institution-wide Cyber exposure estimation. Exhibit 1 below outlines the aforementioned process.

Exhibit 1: Cyber risk quantification process

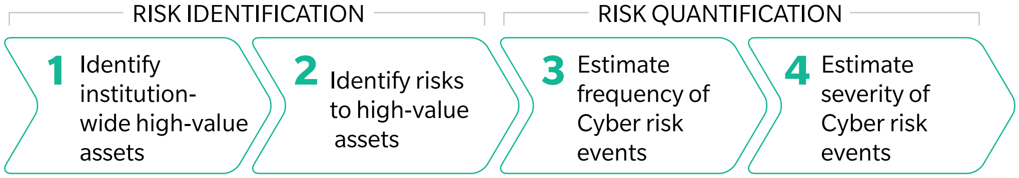

RISK IDENTIFICATION

RISK QUANTIFICATION

1 Identify institution-wide high-value assets

- Develop comprehensive list of Cyber assets
- Determine materiality of identified Cyber assets
- Create institution-wide list of high-value assets

2 Identify risks to high-value assets

- For each high-value asset, identify potential malicious actions
- Assess potential threat vectors and state of existing controls
- Identify most material Cyber risks across institution

3 Estimate frequency of Cyber risk events

- Estimate frequency of Cyber risk events, considering vector of attack, over a given time period
- Supplement analysis with institutional data, simulated attacks, and expert opinion

4 Estimate severity of Cyber risk events

- Use structured scenarios to quantify severity of Cyber risk events based on loss types and loss drivers
- Estimate severity of Cyber risk event scenarios using distributions or point estimates of drivers

STEP ONE: IDENTIFY INSTITUTION-WIDE HIGH-VALUE ASSETS

To begin, business, risk, and information technology personnel should identify assets from all business and functional units (e.g., HR) potentially subject to Cyber attacks. The list should include both digital assets – such as critical data that should be protected or operational services that can be disrupted – and physical assets, including computing hardware and connected infrastructure that can be damaged or destroyed. In the case of digital assets, institutions should create a detailed inventory of not only what kinds of data the institution possesses, but also where and on what servers the data is stored and who has access to such data. Such an inventory of assets is a critical element to understanding Cyber risk exposures across the institution.

Next, still working on the level of business and functional units, Cyber security experts should assess the materiality of each Cyber-relevant asset based on inputs from each business and functional unit. The goal is to identify assets that, if lost or compromised, would lead to significant loss to the institution. In some cases, the asset might have little or no cash value – but loss of the asset might have implications for the ability of the institution to continue to do business. These losses can range from reputational damage (e.g., losing sensitive customer data) to operational interruptions (e.g., inability to process transactions due to system failures). Think, for example, of the Equifax data breach. The data would cost money to replace – but the real materiality of the breach is the effect of such data on the reputation of the institution. The key is to assess each asset from multiple perspectives: How sensitive are the contents of a database? How critical are specific systems and services? What is the potential financial impact of an incursion, or the regulatory or reputational impact? The goal is not so much to assign a hard dollar figure, but rather to rate the relative materiality of each Cyber-relevant asset and develop a list of high-value assets, comprised of the assets that are most material to the institution.

STEP TWO: IDENTIFY RISKS TO HIGH-VALUE ASSETS

Once institution-wide high-value assets are identified, the business or functional unit should develop a list of Cyber risk events by identifying each potential malicious action to which each high-value asset could be subject. The key for the success of the exercise is to consider the most relevant possibilities for potential malicious actions. The unit should be asking who might attack (e.g., a nation state, a professional, an insider, a lone wolf, etc.), what might be the motive for the attack (e.g., financial gain, vandalism, hacktivism, etc.), what form the attack might take (e.g., denial of services, social engineering and phishing, etc.), and what specific assets might be targeted. Exhibit 2 below outlines a framework for the classification of Cyber risks related to high-value assets.

Business, risk, and information technology experts should be engaged to describe the relevant environmental factors for the Cyber risk event. Environmental factors can include the expected strength and expertise of threat agents, the number of potential attack vectors, and the level of controls in place to manage the risk. Controls can consist of preventive measures that reduce the likelihood of a successful attack and of responsive measures that detect and limit the impact of a successful attack.

Exhibit 2: Classification of Cyber risk events

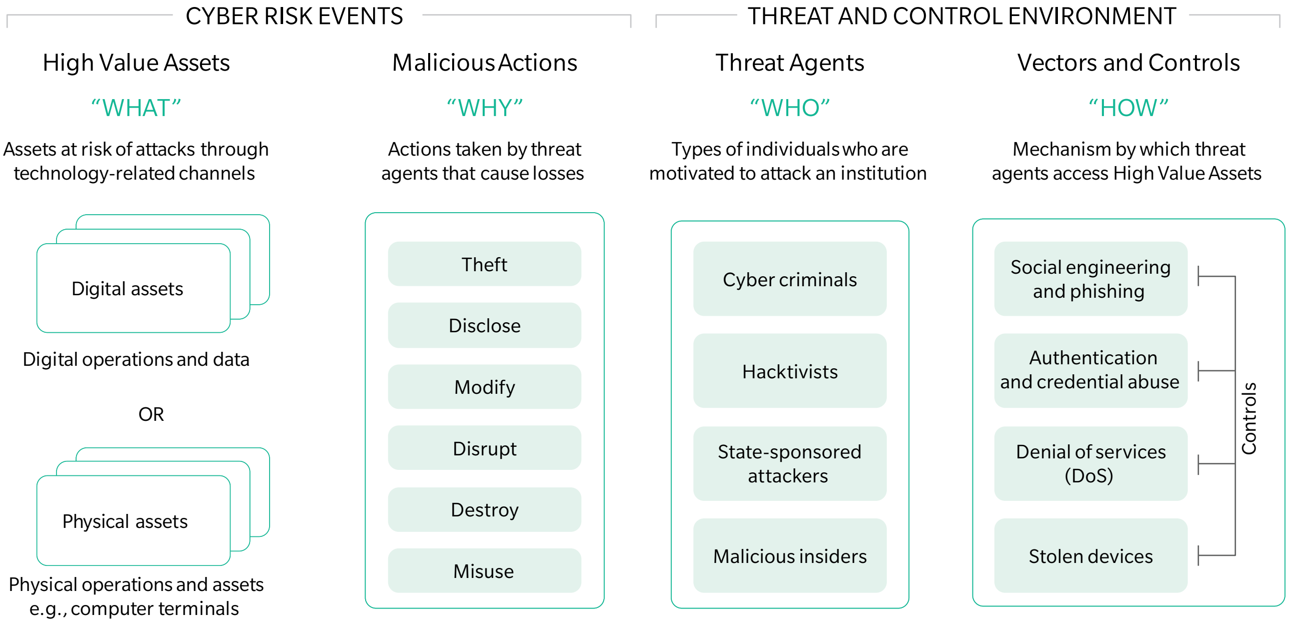
CYBER RISK EVENTS
THREAT AND CONTROL ENVIRONMENT
High Value Assets
"WHAT"
Assets at risk of attacks through technology-related channels
Digital assets
Digital operations and data
OR
Physical assets
Physical operations and assets e.g., computer terminals
Malicious Actions
"WHY"
Actions taken by threat agents that cause losses
Theft
Disclose
Modify
Disrupt
Destroy
Misuse
Threat Agents
"WHO"
Types of individuals who are motivated to attack an institution
Cyber criminals
Hacktivists
State-sponsored attackers
Malicious insiders
Vectors and Controls
"HOW"
Mechanism by which threat agents access High Value Assets
Social engineering and phishing
Authentication and credential abuse
Denial of services (DoS)
Stolen devices
Controls

Finally, the institutions should assign a qualitative risk rating to each Cyber risk event. We suggest assigning ratings based on both inherent risk (potential loss given no controls) and residual risk (potential loss with controls). When assigning qualitative risk ratings to Cyber risk events, institutions should consider both the immediate impact of an attack as well as the longer-term recovery from an attack. For example, a Cyber attack may steal an immaterial amount of money from an institution, but, if the attack is widely publicized, the institution may need years to repair the resulting reputational damage. The result of the qualitative risk rating exercise is a list of the top Cyber risk events to the institution, rated and sorted by materiality and by risk rating. The list can then be used as an input to the quantification process to determine the loss severity of each risk event.

STEP THREE: ESTIMATE FREQUENCY OF CYBER RISK EVENTS

Industry resources exist to help calculate the frequency of various types of Cyber attacks on various sorts of systems in a given time frame. The data allows for a historical view of not only the overall volume of Cyber attacks, but also the volume of attacks and success rates by vector of attack. When analyzing the frequency of Cyber attacks, institutions should consider not only the number of attacks, but also the number of loss-triggering attacks (which are typically a small subset of the total attacks). Furthermore, within these loss-triggering attacks, the magnitude of rare, large loss-triggering attacks drives the tail of the loss distribution, and is therefore most important for institutions to understand, driving the need for better quantification tools of the severity. The data can be a useful starting point, but needs to be used with caution. Cyber crime evolves quickly, and the institution needs to track the changing external environment including threat agents and attack vectors. As necessary, industry data can be supplemented by modeling of institutional data, simulated attacks, and expert opinions.

STEP FOUR: ESTIMATE SEVERITY OF CYBER RISK EVENTS

Given the challenges inherent to traditional quantification approaches commonly used for “financial” risks (e.g., credit, market), we propose to use structured scenarios as a mechanism for quantifying the severity of potential Cyber risk events. These scenarios, which are typically used for the quantification of “hard to quantify” operational risks, will consist of a series of table-top exercises/workshops with key stakeholders from the business, risk, and information technology units. During these sessions, Cyber risk events are broken into basic loss types (including, direct losses, legal fines and fees, regulatory fines and fees, reputational losses, and loss of competitive advantage) and then into more specific loss drivers related directly to the specifics of the risk event. For example, in the event an order management/fulfillment system is disrupted, one underlying driver of losses could be the direct loss of sales that would have occurred during the outage. In addition to direct losses from the Cyber events, institutions should consider whether there will be longer-term recovery costs; for example, institutions may need to invest heavily in marketing and network security following a highly-publicized attack in order to regain customer confidence.

Exhibit 3: Illustration of structured scenario approach

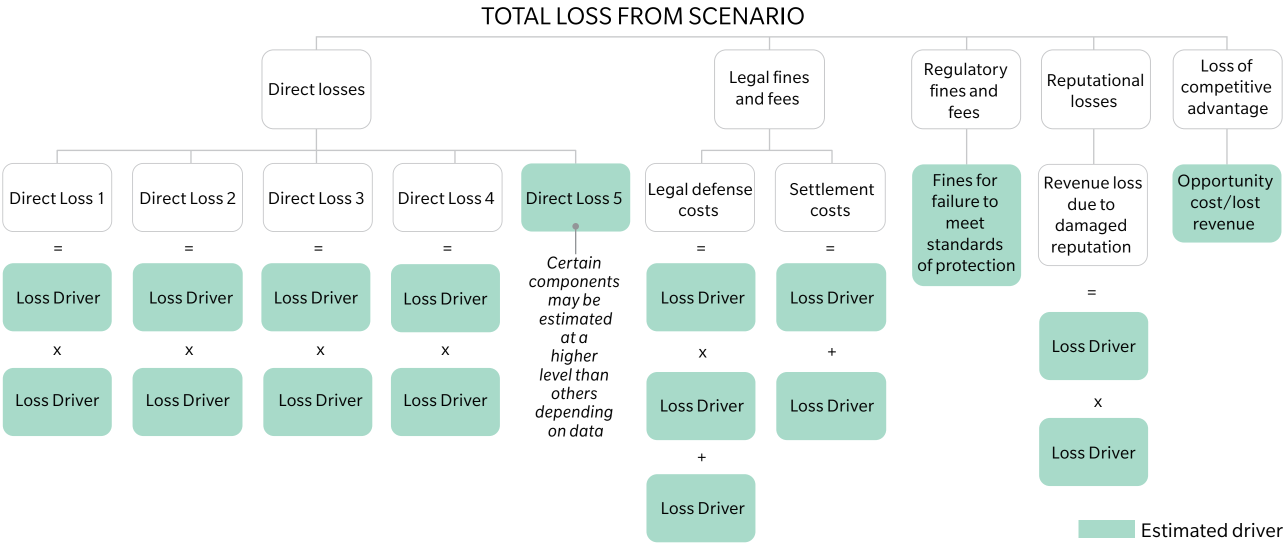
TOTAL LOSS FROM SCENARIO
Direct losses
Legal fines and fees
Regulatory fines and fees
Reputational losses
Loss of competitive advantage
Direct Loss 1
Direct Loss 2
Direct Loss 3
Direct Loss 4
Direct Loss 5
Legal defense costs
Settlement costs
Fines for failure to meet standards of protection
Revenue loss due to damaged reputation
Opportunity cost/lost revenue
=
Loss Driver
x
+
Certain components may be estimated at a higher level than others depending on data
🟩 Estimated driver

Exhibit 3 above represents the decomposition of key Cyber risk events into loss drivers using a tree diagram. Each top node represents a loss type, and the underlying nodes each represent the constituent loss drivers to be estimated, where applicable.

Once the underlying loss drivers have been identified for a Cyber risk event, distributions (or point estimates) can be developed for each underlying driver, based on internal data, external/industry data, or expert judgment. Where insufficient data is available to produce a distribution or where estimates of exact costs are known (for example, hourly costs of employees), a point estimate can be used. However, distributions should be developed where needed to reflect the underlying uncertainty of loss drivers. Expert judgment can be incorporated in the estimates of underlying loss drivers by modifying the distribution or point estimates, but such expert judgment needs to be appropriately justified.

Once distributions and point estimates have been developed for each underlying loss driver, statistical techniques (such as Monte Carlo simulation) can be employed to combine these distributions and point estimates to create loss severity distributions. Exhibit 4 below shows an example output for a sample Cyber risk event (Disruption of Equity Trading Services) for a major financial institution. Loss severity distributions can be aggregated by loss type, and ultimately for the overall scenario.

Exhibit 4: Illustration of output for a sample Cyber risk event

FINAL OUTPUT: DISTRIBUTION OF POTENTIAL LOSSES
SCENARIO: DISRUPTION OF EQUITY TRADING SERVICES

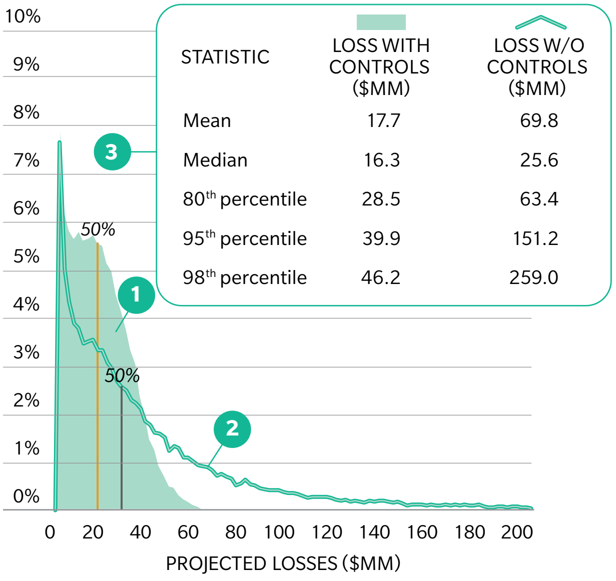

STATISTIC	LOSS WITH CONTROLS ($MM)	LOSS W/O CONTROLS ($MM)
Mean	17.7	69.8
Median	16.3	25.6
80th percentile	28.5	63.4
95th percentile	39.9	151.2
98th percentile	46.2	259.0

OUTPUT COMMENTARY

1. A distribution of potential losses under a particular cyber risk scenario can be generated
2. Modifications and adjustments made to underlying drivers can be applied to the model, which results in shifted distributions
   - These changes can reflect the impact of changing assumptions such as increased controls, or control failures
   - Comparisons can be made to understand sensitivities to contributing drivers
3. The percentile outputs can be used to represent both the expected severity of a typical occurrence as well as severities under less likely, but more extreme outcomes
   - The mean and median represent measures of the expected severity of the loss event
   - Higher percentiles represent "tail-events," which can be used to represent worst-case outcomes

MANAGING RISK FOR TODAY AND TOMORROW

The objective of risk quantification, of course, is to use the output to allow for more informed business decisions and improved Cyber risk management – to ensure that the organization is operating consistently within stated risk appetite, to identify priorities for developing response strategies, to help make decisions about preventive measures and controls, and to inform where insurance might be appropriate to manage losses. Exhibit 5 below summarizes some key business applications of our proposed approach.

Exhibit 5: Key business applications of Cyber Risk identification and quantification

RISK MANAGEMENT

Better understand Cyber Risk exposure and the underlying drivers of the losses, and improve response to attacks

INVESTMENTS

Prioritize investments across the Cyber Risk mitigation spectrum and relative to competing investment demands

INSURANCE

Determine Cyber coverage strategy and the nature/extent of premiums

EXECUTIVE OVERSIGHT

Understand Cyber Risk exposure status, trends/outlook and impact of investments over time

By quantifying the exposure to Cyber risk, organizations will be better positioned to take all those steps. Significantly, by converting qualitative concerns from Boards and senior management into dollar amounts, an institution will be able to integrate Cyber risk management more fully into the overall risk management strategy – which is the ultimate goal.

We believe the quantification exercise is valuable because of the numerous uncertainties surrounding Cyber risk. The structured inquiry we propose is designed not just to produce an improved (though, alas, still imperfect) estimate of Cyber risk exposure today, but to enhance the understanding of risk managers and business and functional units alike, and to launch the sort of feedback loop that over the course of multiple iterations will give institutions deeper, more reliable, and more actionable understanding of the risks faced in an increasingly digital world. Such understanding is vital today and will only grow in importance in a riskier and more complex future.

Oliver Wyman is a global leader in management consulting that combines deep industry knowledge with specialized expertise in strategy, operations, risk management, and organization transformation.

For more information please contact the marketing department by email at info-FS@oliverwyman.com or by phone at one of the following locations:

AMERICAS
+1 212 541 8100

EMEA
+44 20 7333 8333

ASIA PACIFIC
+65 6510 9700

ABOUT THE AUTHORS

RAMY FARHA
Partner in the Finance & Risk and Public Policy Practices
ramy.farha@oliverwyman.com

EVAN SEKERIS
Partner in the Finance & Risk and Public Policy Practices
evan.sekeris@oliverwyman.com

JERRY WU
Engagement Manager in the Finance & Risk and Public Policy Practices
jerry.wu@oliverwyman.com

www.oliverwyman.com

The information and opinions in this report were prepared by Oliver Wyman. This report is not investment advice and should not be relied on for such advice or as a substitute for consultation with professional accountants, tax, legal or financial advisors. Oliver Wyman has made every effort to use reliable, up-to-date and comprehensive information and analysis, but all information is provided without warranty of any kind, express or implied. Oliver Wyman disclaims any responsibility to update the information or conclusions in this report. Oliver Wyman accepts no liability for any loss arising from any action taken or refrained from as a result of information contained in this report or any reports or sources of information referred to herein, or for any consequential, special or similar damages even if advised of the possibility of such damages. The report is not an offer to buy or sell securities or a solicitation of an offer to buy or sell securities. This report may not be sold without the written consent of Oliver Wyman.