FORESIGHT
RESILIENCE
STRATEGIES

Increasing Cyber Resilience through a Risk-Based Cybersecurity Program

Building a Bridge from Qualitative to Quantitative Measurement

Executive Summary

Organizations know they need to take a risk-based approach to cybersecurity but struggle with what such an approach would look like.

Large companies with sophisticated cybersecurity programs frequently suffer cyber incidents, so smaller organizations conclude they are even more vulnerable and need to increase their cyber spending. But by how much?

Finding a common language to discuss cybersecurity is also a challenge. In addition to translating from technical language to the language of business objectives, a solution should clearly illuminate which strategies, controls, and tools will help deliver measurable improvements and should quantify the return on cybersecurity investments for business leaders in risk-based terms. For the information security team, the ability to communicate both the risks and risk mitigations in terms that the business leadership can understand is invaluable.

Among the possible approaches to building a risk-based cyber program, the NIST Cybersecurity Framework stands out as the most widely-cited approach in use today, and the discussion below is aligned to this framework. This discussion should feel familiar for any other established risk-based framework.

Beyond describing how a risk-based approach works, we will cover in some depth the limitations and pitfalls of qualitative implementations that rely on ordinal measures such as “low,” “medium,” and “high,” commonly found in risk matrices, or heat maps. We do not dismiss the value of such tools. Beginning with a qualitative method can offer significant benefits. Rather, the goal of this white paper is to propose orienting such a qualitative approach so that it seamlessly supports a move towards quantification as the next logical step.

In this paper, we hope to make the case for an approach that ensures that initial efforts necessary to producing a qualitative risk assessment can be used as a solid foundation for a robust quantified risk assessment by:

- Starting with information assets
- Evaluating different types of loss
- Estimating loss in quantitative terms

Foresight Resilience Strategies helps clients develop and improve their risk management programs, whether by adding risk quantification to an existing qualitative approach or by building a quantitative program from scratch.

Introduction: Tools for Building a Risk-Based Cyber Program

The foundational document in cybersecurity strategy today is the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF).[1] Since its original release in 2014 and even more since the release of Version 1.1 in 2018,[2] risk management has become the watchword of cybersecurity practice. The federal government requires CSF usage by agencies in developing their cybersecurity plans.[3] Among critical infrastructure companies –– its target audience –– the CSF is widely respected and considered a best practice.[4] Though designed for critical infrastructure, "organizations worldwide have found [the CSF] to be a prioritized, flexible, repeatable and cost-effective approach for managing cyber risk."[5]

The CSF charges users to view cybersecurity programs through the lens, language, and methodology of enterprise risk management (ERM). When measuring the risk, that means using the tools of actuarial science to ***model the potential loss as a product of the impact of a scenario and its likelihood.*** The analytical results generated by an ERM model provide a means to prioritize activities and manage risk by choosing whether to avoid, transfer, mitigate, or accept the risk.

The CSF builds on this idea to provide users with a way to organize cybersecurity activities into five functions: Identify, Protect, Detect, Respond, and Recover. These functions are in turn divided into over 100 categories and a larger number of sub-categories, which contain the specific activities and even the specific tasks that make up a cybersecurity program. This hierarchical and

[1] *NIST Framework for Improving Critical Infrastructure Cyber Security*, released February 2014, revised February 2018 (accessible at https://www.nist.gov/cyberframework).

[2] Version 1.1 was released in February 2018 (accessible at https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.04162018.pdf).

[3] Executive Order 13800.

[4] *NIST Cybersecurity Framework Adoption on the Rise*, Tenable (https://www.tenable.com/blog/nist-cybersecurity-framework-adoption-on-the-rise).

[5] 75% of Organizations Are at Significant Risk of Cyber Incidents, *RSA* (https://www.rsa.com/en-us/company/news/rsa-research-75-of-organizations-are-at-significant-risk-of-cyber-incidents).

organizational approach provides a way of thinking about cybersecurity activities but is not at the center of introducing risk as the main way to assess a cybersecurity program.

The CSF –– especially the Framework Core, the basic document that NIST maintains on its website –– intentionally remains fairly abstract and general in order to cater to an extremely wide range of organizations.[6] NIST states that it seeks to make the CSF more useful through the publication of Profiles, one of the three elements considered part of CSF itself, and other resources (including several provided by NIST), all of which provide a roadmap to reduce the CSF to practice.

Getting the most out of the CSF

An organization trying to implement a risk-based approach to cybersecurity will need to look beyond the four corners of the CSF in order to find tools to estimate information risk in their organization. This raises two immediate problems. First, the most important initial activity in assessing a cybersecurity program –– something that would be viewed as a precursor to implementing controls –– is a key cybersecurity control itself: determining what information the organization holds, what network assets that information is stored on, and what users have access to different information assets. The immediate question is then whether an organization has to begin the sometimes arduous process of choosing and beginning to adopt a control system to implement before it has assessed what its needs are.

Instead, organizations often start by creating a risk register. Listing risks feels like less of a commitment because it does not involve selecting a control system. ***In fact, creating the range of scenarios that a complete risk register should contain means devoting significant resources to the process of imagining all possible negative outcomes.*** In addition, involving a large number of people –– necessary because of the desire for a comprehensive result –– means that the register will likely include things not actually risks:[7] vulnerabilities, threats, actors, and conditions frequently find their way onto a risk register. Because the resulting product is not aligned to compare risks with one another, the tool can be flawed from the outset.

[6] That primary document also contains the second of the three primary elements of the CSF, the Implementation Tiers. Those do not form a significant part of this discussion.

[7] The word risk is often used in imprecise ways that undermine the effort to talk clearly about the solutions. Here we are using risk to refer to uncertain, significant, and specific future outcomes and distinguish risk from the conditions, actors, and circumstances that lead to those outcomes.

The second problem that arises when users go beyond the four corners of the CSF is that the essential qualitative tool of risk management, ***the heat map or risk matrix, tends to provide less actionable outputs than required when making business decisions.*** A heat map involves identifying the impact and likelihood of a risk on a simple, qualitative scale, usually high/medium/low or 1 through 5. Because those estimates of impact and likelihood do not account for the possibility that either may fluctuate, heat maps tend to misstate the accuracy of the estimate of impact and likelihood of a risk by overstating the precision of that estimate. In addition, there is not a clear relationship between the low, medium, and high values in the qualitative scale.[8]

For example, if a doctor's office rates the loss of patient health information a high impact event (eminently reasonable on its face), how does that account for the relatively minor impact of a single patient record being faxed to the wrong outpatient surgery center. Typically, this small mistake can be addressed by calling the center to ask that they destroy the fax and reporting to the U.S. Department of Health and Human Services for information purposes only at the end of the calendar year.[9] There is a range of potential impacts that a simple rating of "high" does not capture.

A heat map also struggles to inform decisions about spending, either from a forward-looking budget perspective or from a retrospective review of previous expenditures. Typically, regions of the heat map translate to the colors of a traffic signal: red risks are high likelihood and high impact and should be mitigated first. Green risks are low impact and low likelihood and can be considered under control. That leaves yellow risks. ***How much should an organization spend to mitigate a yellow risk? And should it spend the same amount for every one?***

Adopting a Quantitative Approach

The CSF does not require the use of qualitative estimates and heat maps; and a quantitative approach to information risk is entirely consistent with the CSF.[10]

[8] Do two lows equal a medium and two mediums equal a high? That is almost never defined.

[9] HHS requires prompt notification of breaches that involve over 500 records and publishes notifications of all of those breaches. Breaches of fewer than 500 records are not published on their website. https://www.hhs.gov/hipaa/for-professionals/breach-notification/breach-reporting/index.html

[10] Given that the original charge to NIST wanted the resulting framework to offer a cost-effectiveness way to manage cyber risk, it could be argued that a quantitative approach is preferred. *See* Executive Order 13686, §7(b).

The CSF also does not go into detail on the limitations of a qualitative approach. In fact, organizations that have adopted such an approach have taken an important step toward understanding cybersecurity as a risk-based discipline. So long as the limitations are understood, it is certainly better to begin with a qualitative approach than not to use a risk-oriented approach at all.

While a qualitative risk model provides useful benefits to organizations initially, anyone beginning the cyber risk management journey in 2019 may want to simply begin with activities that will support a quantitative approach, either immediately or eventually. Injecting certain quantifiable characteristics into a primarily qualitative analytic model will ensure that when a more sophisticated approach is desired, the foundation for such an approach will already be in place.

Conclusions and Recommendations

There are important techniques and considerations one should embrace so that the efforts involved in a qualitative risk assessment will transfer easily to a quantified one when the time comes. Here are the important considerations.

Don't start with scenarios: Rather than working from a variety of possible outcomes, start with the information assets held and consider potential losses related to those information assets regardless of the reasons or narratives around those losses. This will support probability-based analysis later.

Build up multiple dimensions: All losses of control are not the same. The impact of a loss of confidentiality might be more important than a loss of integrity or availability for some information (e.g., medical records). The opposite is true for other information.

Recognize different types of loss: In addition, loss of control over information can result in many different types of loss––lost sales, replacement costs, reputational damage––consider losses of all these different types when assessing the magnitude of impact and likelihood. That will help to determine what controls to focus on for any particular information asset.

Consider bringing quantification in early: Rather than using low/medium/high for impact, why not make several bands identifying ranges of losses in dollar values. Similarly, likelihood is most meaningfully expressed as a probability of an event occurring in a specified time period.

Using this as the initial approach, the results will drive better decision making. Rather than having a red/yellow/green understanding of risk, the analysis will point toward what information control and type of loss should be of concern. By using initial cost estimates when assessing the impact, the appropriate controls will be easier to identify.

The next step is to get more specific about the elicitation of the ranges of impact and to find a source for better data on the frequency of incidents. The results of a more sophisticated analysis will allow cost-effectiveness calculations for the current cybersecurity program and provide return on investment forecasting for future spending.

Schedule a 30-minute consultation with Foresight today to learn how we can help you improve your assessment to build a quantitative risk-based cyber program that communicates risk effectively across your organization and provides analytical results in financial terms.