Board Oversight of Cyber Risk

BASELINE DIAGNOSTIC GUIDE

A simple five-step guide to help CEOs and boards of directors gauge their organization's ability to accurately measure and prioritize cyber risk

By Jack Jones
Chairman, FAIR Institute

Overview

Board directors and senior executives are obligated to govern their organizations' cyber risk management efforts. To-date, the information they've had to rely on in this effort has focused primarily on industry benchmarks, compliance metrics, and a bewildering array of security-related data. Unfortunately, these types of reports fail to capture some key elements that are crucial to risk management success.

The premise underlying this document is that effective risk management can only occur when an organization can effectively prioritize its efforts. This is particularly true given the complex and dynamic nature of the cybersecurity landscape, and the inescapable reality of limited resources. The ability to prioritize effectively is predicated on being able to reliably and clearly compare and prioritize the challenges facing an organization, and the available solutions.

If you examine the risk management standards and benchmark frameworks commonly used in cybersecurity (e.g., NIST CSF, ISO2700x, etc.) you will notice that although they all call for risk measurements to be performed, none of them provide specific guidance on how to perform risk measurements. And if you investigate how risk is actually being measured in organizations, you'll find little or no rigor, and the measurements themselves rarely stand up to critical examination. Furthermore, the typical high-medium-low characterization of risk does not support the kind of economical trade-off decisions that are necessary in order to allocate resources cost-effectively.

The purpose of this document is to provide a simple, high-level diagnostic tool that boards, and senior executives can use to gauge their organization's ability to accurately prioritize its cyber risks and apply cyber risk resources cost-effectively. It also provides a set of recommended actions to correct any identified deficiencies.

Answers to the questions described in this document should most likely come from the organization's CISO[1]. Note that in the process of developing this tool, these questions were posed to a number of CISOs. None of the CISOs argued against the underlying premise or the questions themselves, and none of them felt they could answer any of these questions satisfactorily. This unfortunate state exists because the cybersecurity field has historically focused on technical or regulatory compliance rather than economical, business-driven decision-making based on meaningful risk measurement.

Question 1: Do we know what our crown jewels are and where they're located?

Diagnostic purpose: To understand whether the organization is able to appropriately protect the assets that represent the most value and/or liability to the organization.

Most likely answer(s): "No" — to both parts of the question, or "Yes" and "No" respectively. For this second answer, it is often the case that some attempt has been made to define an organization's crown jewels but very often these definitions are created in a vacuum and don't effectively represent the business perspective.

Background: Technology asset management is rarely a priority for organizations, which means that very few organizations today have (or at least reliably execute) the policies, processes, or technologies that allow them to reliably know which assets qualify as "crown jewels", or where those key assets exist. In

[1] The term CISO (Chief Information Security Officer) is used throughout this document, however organizations may have other titles for personnel with the same responsibilities, such as VP of Information Security, VP of Technology Risk, etc.

fact, most organizations have not taken the time to explicitly define what constitutes their crown jewels. As a result, an organization’s ability to protect itself from truly significant loss must be suspect.

Because this concern is so rarely focused on, if a CISO answers that the organization does, in fact, know what and where its crown jewels are, the board should seek validation by a third party.

Examples of crown jewels include, but aren’t limited to:

- Key intellectual property
- Sensitive corporate strategies and/or plans
- Large stores of sensitive consumer/customer information
- Critical business processes, applications and the technologies that support them
- 3rd parties that provide critical services and/or that have access to large volumes of sensitive information

The more explicitly these are defined, the easier it is to find, manage, and appropriately protect them.

Note that it can be tempting to think of every sensitive customer or employee record as a crown jewel, but this is neither practical nor achievable from a risk management perspective. A critical part of effectively managing any problem is the ability to focus, and the saying; “*When everything is a priority, nothing is*.” applies here as well.

Recommended actions: If the answer to the first part of this question is “no”, then the assignment should be for the organization to clearly define the criteria for “crown jewels” and then locate and keep track of them. If crown jewels have been clearly defined but their locations are uncertain, then the task is to implement or strengthen the policies, processes, and technologies to identify and keep track of them.

Question 2: What are our top ten cyber risks?

Diagnostic purpose: To evaluate whether your organization has a clear and accurate understanding of what constitutes a cyber risk, which is a fundamental prerequisite to accurate measurement and effective prioritization.

Most likely answer: Most organizations will answer this question with a list that primarily consists of risk landscape components, but that are not in-and-of-themselves risks. The following is an example of what is commonly (and incorrectly) presented as “top risks”:

- Cyber criminals
- Disgruntled insiders
- The cloud
- Sensitive customer information
- Reputation
- Weak passwords
- Patching (or unpatched systems)
- Lack of awareness
- Phishing
- Ransomware

If you are given a list like this, there is very little chance that your organization can measure and prioritize cyber risk effectively.

Background: Risk is measured in terms of the likelihood and magnitude of loss, which can only be applied to loss event scenarios. These scenarios require at least two elements — an asset and a threat to that asset, for example:

- A cyber-criminal compromises sensitive customer information using a phishing attack
- A disgruntled employee or contractor takes down critical systems using ransomware

The two examples above include a third element, which is the method (e.g., phishing, ransomware) used by the threat actor (e.g., cyber-criminal, disgruntled employee). This provides additional specificity that improves both risk measurement precision and the ability to identify cost-effective control opportunities.

Although elements from the former (errant) list were combined to define these two loss event examples, the items in the former list are not in-and-of themselves loss event scenarios, which means likelihood and magnitude of loss cannot logically or reliably be assigned to them individually. Unfortunately, most organizations don't recognize this distinction and therefore inaccurately assign likelihood and impact ratings to the items in their risk register. This is one of the most significant contributors to inaccurate risk measurement, ineffective prioritization, and inefficient application of risk management resources.

If your organization's list of top risks and/or its risk register contain anything other than loss event scenarios (i.e., if it is made up of items like those in the former list), then you can be certain that your organization is unable to reliably prioritize its risk management efforts.

Recommended actions: Insist that the organization review and reconcile its risk register and/or list of top risks to ensure that they contain actual risks, according to a standard risk taxonomy such as Factor Analysis of Information Risk (FAIR). This often requires that personnel responsible for this effort recalibrate how they think about risk in general, and how they use the term "risk". Once any necessary adjustments have been completed to the risk register/list, risk measurement values can be reevaluated and reassigned.

Question 3: How much loss exposure (in economic terms) does the top cyber risk represent?

Diagnostic purpose: Contrary to what you might imagine, this question isn't intended to just provide you with an economic expression of cyber risk loss exposure. In fact, its main purpose is to help you gauge the CISO's understanding of modern risk measurement methods as well as their receptiveness to quantitative methods.

Most likely answers: Many CISOs will answer this question in one of the following ways:

- "Cyber risk can't be measured in economic terms."
- "I don't know but will look into it and get back to you."

Background: There has been a long-held belief in the cybersecurity profession that cyber risk is somehow different from other forms of risk and can't be measured in economic terms. Furthermore, methods for effectively measuring cyber risk in economic terms are relatively new, with adoption growing only within the past few years. As a result, it should not be too surprising if your CISO isn't familiar with these methods or believes it can't be done.

A CISO's familiarity with these methods can easily be corrected by pointing them to the resources listed at the end of this document, and by having the CISO speak with organizations that have been successful in measuring cyber risk in economic terms. What is often more challenging is for CISOs to overcome preconceptions and biases regarding quantitative methods. It is not uncommon to encounter significant resistance on this subject.

Recommended actions: If your CISO is unfamiliar with methods for measuring cyber risk in economic terms, they should be directed to become familiar with them, and to report back on the subject. If they come back still claiming that it can't be done, then serious consideration should be given to finding a new CISO, as their response suggests an unwillingness to evolve the organization's capabilities.

Question 4: Who is allowed to measure cyber risk in the organization?

Diagnostic purpose: To gauge whether the organization recognizes that reliable risk measurement requires certain skills.

Most likely answer: “Anyone"

Background: In most organizations, anyone within the risk, audit, security, or technology organizations is allowed to rate/measure cyber risks (or, more often as discussed in question 2 above, non-risks). The problem is that reliable risk measurement requires certain fundamental skills that are often not present in the personnel who do risk ratings/measurements. These skills include:

- Critical thinking. The cyber profession's reliance on compliance checklists and other superficial frameworks tends to atrophy the ability and willingness to think critically about the complex nature of cyber-related problems. However, reliable measurement of cyber risk requires a strong ability to deal with complexity and uncertainty.
- Being comfortable with numbers. Measurement (even when those measurements are expressed qualitatively) involves data and numbers, so it is logical to expect that personnel involved in risk measurement should be comfortable working with numbers.
- Understanding basic probability principles. Risk is inherently a question of probabilities; therefore, it stands to reason that personnel engaged in risk measurement have a basic understanding of these principles. However, as with numeracy, the skills here do not have to be especially deep for most day-to-day risk measurement.
- Relying on a standard risk model. Every risk measurement involves a model. These models range from the commonly unexamined and uncalibrated mental models of your average professional, to formally defined and vetted models like FAIR. Personnel who are authorized to measure cyber risk should be required to understand and leverage a clearly defined and vetted model.
- Being a calibrated estimator. Most people are extremely poor at estimation, which means that when someone waves their wet finger in the air to give a gut-driven estimate, the odds of that estimate being accurate is low. Fortunately, most people can learn how to estimate well in just a few hours of training using techniques that are well-established.

The bottom line is that being stellar at auditing, security architecture, compliance, or some other risk/security role does not automatically mean someone has the ability to measure risk reliably.

Recommended actions: Consider establishing a policy that all cyber risk measurements must be performed (or at least reviewed) by personnel who meet the above criteria. This is especially crucial for identifying and measuring top risks, and for cost-benefit analysis on major risk management investments.

Regardless of how your organization assigns risk measurement responsibilities, you should make it clear that risk measurement is considered a crucial role, and anyone performing that responsibility must be qualified.

Question 5: What is the prevailing root cause behind execution failures and deficient controls?

Diagnostic purpose: To understand whether the organization is able to identify and treat common/systemic causes of execution failures (which drives most non-compliance, security failures, and poor performance).

Most likely answer: "Awareness"

Background: Most organizations fight the same cyber risk management battles over and over (e.g., persistent access management failures, missed remediation deadlines on audit findings, missed deadlines on applying critical patches, unauthorized systems being installed on the network, etc.). This not only means that organizations repeatedly experience additional and perhaps unacceptable levels of loss exposure when/where these problems exist, but they also waste resources applying band-aids rather than meaningful solutions.

This is invariably a function of one or both of the following:

- Failure to perform root cause analyses at all
- Performing proximate cause analysis rather than root cause analysis

There is a prevailing belief that lack of policy awareness is the largest contributing factor to non-compliant conditions. In fact, root cause analyses performed in a number of organizations suggest something different. Without exception, deeper root cause analyses performed in large enterprises identified that 70% to 80% of security failures occurred when the responsible parties were aware of what was expected of them and had the necessary skills and resources. The responsible parties simply chose not to comply because they considered deadlines and/or budget constraints to be more important, or compliance was simply considered too inconvenient.

These choices usually carry no consequences for the responsible parties for one or both of the following reasons:

- The policies and standards they're expected to comply with are aspirational and not well-aligned with the organization's risk appetite.
- Management does not truly understand the loss exposure implications of non-compliance, and therefore doesn't incentivize it on even footing with other business imperatives.

Note that poor risk measurement plays an important role in both of these reasons.

Another important point to keep in mind is that root causes can often be outside of the CISO's authority to fix, which makes senior executive support imperative.

Recommended actions: Steps for remedying this concern include:

- Improving the organization's ability to measure and report risk accurately (see earlier recommendations).
- Reviewing and reconciling policies to align with the organization's risk appetite.

- Establishing policies and processes for strong root cause analysis.
- Holding people accountable for willful noncompliance.

Summary

Today's risk management standards and compliance frameworks, as well as common maturity models, can be useful risk management tools. They do not, however, focus on the fundamental need to measure risk accurately so that risk management decisions can be well-informed and business-aligned. Nor do they address the need to identify and treat the root causes behind execution failures.

By examining your organization through the lens provided by the five questions discussed above, you can better understand your organization's ability to prioritize its risk management efforts effectively and treat the root causes of execution failures.

Resolving any deficiencies identified using these questions should position your organizations to be much more cost-effective in their risk management efforts, and therefore less prone to major cyber-related loss events. It also should lay the foundation for your CISO to communicate with senior executives and the board in economic terms that are more meaningful and accurate, and which enables you to govern more effectively.

Resources

Correcting any problems identified through these questions often requires some amount of training and recalibration for personnel involved in the effort. Sources of additional information, training, services, and tools, both free and commercial, include:

Websites:

- The FAIR Institute (www.fairinstitute.org)
- The Open Group (www.opengroup.org)
- RiskLens (www.risklens.com)

FAIR and FAIR-based Training Courses:

- FAIR Analysis Fundamentals (www.risklens-academy.myshopify.com)
- FAIR Analyst Learning Path (www.risklens-academy.myshopify.com)
- Resolve: Board and Executive Cybersecurity Training (www.cybervista.net/resolve)

Reference books:

- Measuring and Managing Information Risk: A FAIR Approach (the Amazon book store)
- How to Measure Anything in Cybersecurity Risk (the Amazon book store)