WHITE PAPER

CYBER RISK MANAGEMENT MATURITY

Written for the FAIR Institute
Cyber Risk Workgroup
March 2017

Jack Jones
EVP Research & Development, RiskLens

Laying the foundation for maturity

You can't swing the proverbial "dead cat" without hitting yet another cyber risk maturity framework (e.g., NIST CSF, FFIEC CAT, and many proprietary and home-grown approaches). Almost without exception, these frameworks are made up of a patchwork of controls, processes, and policies believed to represent "good practices". The underlying premise is that the more completely aligned an organization is to these practices, the more "mature" the organization is. Although this is logical on the surface, these frameworks do not represent maturity in any real sense.

This document will describe a more fundamental approach to defining and evaluating cyber risk management maturity. In doing so, it will also provide a high-level framework for evaluating an organization's cyber risk management maturity level across various foundational elements. This approach to evaluating maturity will align with some elements in other frameworks, as well as fill gaps within those other frameworks.

An organization can use the information in this document to better understand its current state of cyber risk management maturity and develop a road map for evolving that maturity.

Underlying principles

There are just a handful of foundational principles that underlie the approach described in this document. They are:

- Cyber risk is just one of the many concerns that leadership has to apply its limited resources to, the others being: opportunities, operational needs, and other forms of risk.
- The objective of cyber risk-related policies, processes, practices, personnel, and technologies is to help ensure that an organization's exposure to cyber-related harm is maintained at a level that is acceptable to its leadership.
- Achieving and maintaining an acceptable level of cyber-related risk can only occur if decision-makers are well-informed about their risk landscape.
- The benefits of well-informed decision-making can only be realized if execution against those decisions is reliable.

Given the above, <u>a mature cyber risk management program is one that can achieve and maintain an appropriate level of cyber-related risk</u>, which can only be reached through well-informed decision-making and reliable execution.

Some readers may be concerned with the fact that common cyber risk controls like authentication, patching, logging, etc. aren't included in this model. These readers should keep in mind that any organization making <u>well-informed</u> choices about cyber risk will implement appropriate controls of those types. In other words, appropriate control choices are expected to follow from well-informed decisions.

The risk landscape as a system

The diagram below provides a visual description of the risk landscape as a system, and the role of decisions and execution in managing risk. The most important point to take away from this diagram is the feedback loop that provides decision-makers with intelligence regarding their risk landscape. The better this feedback loop is operating, the better able decision-makers will be to make appropriate risk management choices. Note, too, that this feedback loop includes not just information about risk (threats, assets, controls, and impact) but also information regarding the efficacy of risk management practices (decision-making and execution).

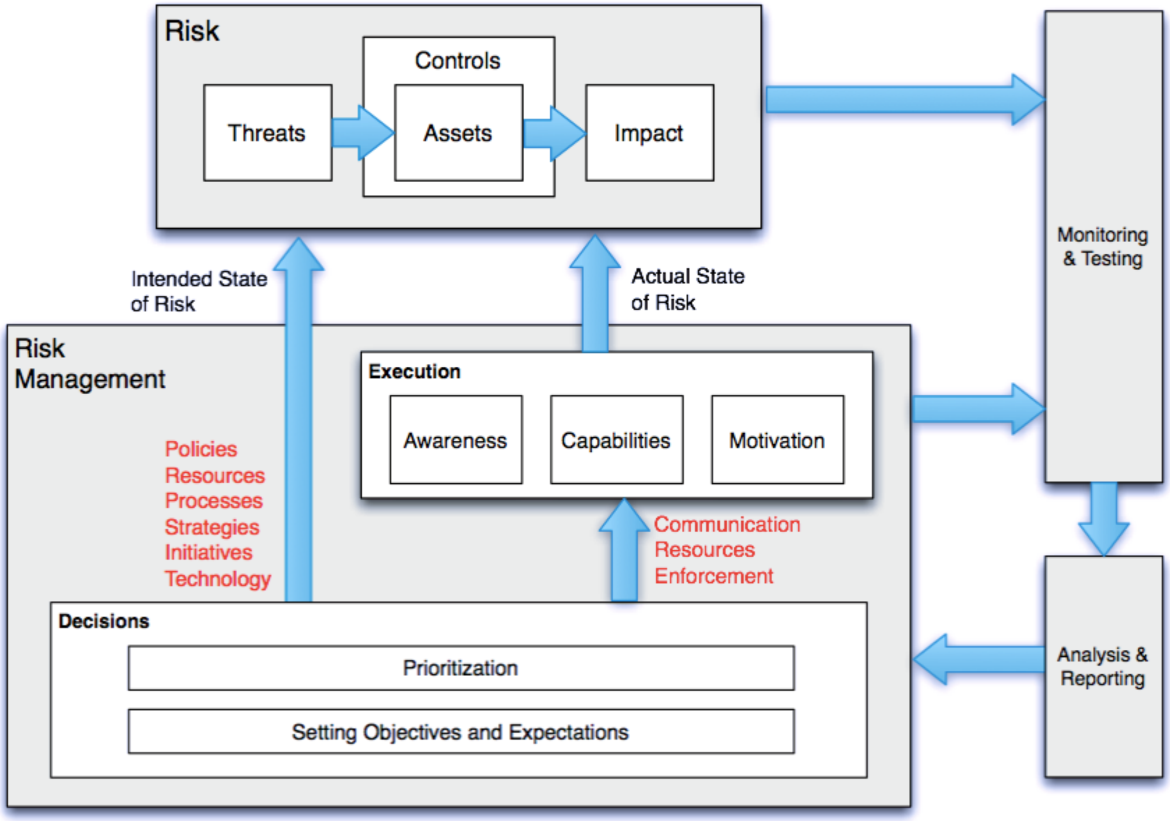

The next section discusses the factors that support well-informed decisions and reliable execution, as these are the elements organizations can evaluate to determine their level of cyber risk management maturity. The last section of the document will provide an example maturity model you can use to evaluate your organization’s cyber risk management maturity.

Well-Informed decision-making

Risk management begins with decision-making, typically in the form of policies that set expectations for the risk management program. Of course, decisions extend well beyond policies. Some of these include:

- Setting risk management objectives — e.g., defining risk appetite
- Strategies for achieving risk management objectives
- Plans for carrying out those strategies
- Resource allocations
- Technology standards
- Approvals to policy exceptions
- Responses to cyber-related events
- Prioritization of activities, both strategic and tactical
- Adjustments to strategies, plans, activities, etc. as the risk landscape evolves

All of these decisions will affect an organization's risk profile, either positively or negatively. They also affect the organization's ability to pursue opportunities, meet operational needs, and deal with other forms of risk. As a result, it is imperative that these decisions be as well-informed as possible given available resources.

Risk terminology

All decisions are choices, which implies making comparisons between choices, which implies the ability to measure the relevance/value of the choices on a common scale. Furthermore, consistent/reliable measurement is predicated on clear definitions of the things being measured. Therefore, no organization can claim to be mature if personnel involved in risk management operate from inconsistent and/or unclear risk-related terminology.

Unfortunately, most organizations today have not adopted nor enforce a clear and consistent set of risk-related terminology. This profoundly affects the quality and consistency of risk measurements and communications, which prohibits well-informed decision-making.

NOTE: In terms of developing a roadmap toward risk management maturity, this is where you start the journey. Fix this first. By the way, this is one of the first and most important value propositions organizations realize from adopting FAIR.

Risk landscape visibility

You can’t manage what you aren’t aware of.

Assets

An organization must have reasonably accurate information about its assets at risk. This involves more than just whether an asset exists in the risk landscape, but also the value/liability characteristics of the asset, which largely determine the loss implications of any adverse event. The challenge here is that many organizations have highly dynamic asset populations (systems, network connections, applications, etc.) and don’t have mature processes for keeping track of these assets.

Threats

One of the most challenging aspects of cyber risk management is a threat landscape that is highly dynamic and made up of threat communities that are both external and internal to the organization. Furthermore, threat intelligence includes both a tactical and strategic point of view — i.e., what’s going on right now and is likely to happen in the near future, versus how the threat landscape is evolving. As a result, making well-informed risk-based decisions requires that an organization have the highest quality threat intelligence it can afford.

The maturity of threat intelligence can be thought of in three dimensions — timeliness, reliability and scope. The importance of timeliness and reliability is relatively obvious. Scope, however, is less obvious but can be broken down as:

- At the lowest level of relevance to an organization is threat intelligence that is highly generalized and is not specific to the organization’s industry or to the organization itself.
- At the next level of relevance is threat intelligence that is specific to an organization’s industry.
- The most relevant threat intelligence is that which is specific to the organization itself.

Due to the highly-specialized nature of the threat landscape, all but the largest organizations usually must engage external expertise (threat intelligence providers) to some extent in order to ensure high quality threat intelligence.

Controls

Accurately evaluating risk requires good information regarding the condition of the controls that are intended to protect assets. Due to limited resources, however, it is infeasible to maintain high levels of visibility into all of the controls across all of an organization's assets. Consequently, organizations should strive to maintain a level of controls visibility that is based on three factors:

- The value/liability characteristics of an asset (i.e., the loss implications)
- The threat activity an asset faces, which is often driven by a combination of its perceived value to threat communities and its exposure to those communities
- The frequency of changes to the asset that could affect control conditions (e.g., how often the asset is subject to configuration changes, etc.)

Prioritizing control visibility efforts like auditing, scanning, and testing based on these parameters can help to ensure that organizations achieve cost-effective visibility.

Execution

Variance in control conditions occurs primarily due to poor execution, and these variant conditions contribute to virtually every cyber-related loss event organizations experience. As a result, it is vital that organizations recognize and correct the conditions that contribute to poor execution. Many (most?) organizations however, suffer from "risk management Groundhog Day" — i.e., they experience the same execution failures repeatedly. Examples of these include things like repeat audit findings, inconsistent access management compliance, system misconfigurations, etc.

Performing root cause analysis on non-compliant (variant) control conditions is the key to recognizing and systemically correcting variance. Although some organizations claim to perform root cause analysis on non-compliant conditions, the vast majority of these analyses are superficial and only identify proximate causes. Furthermore, it is exceptionally rare to encounter an organization that evaluates root causes systemically across their portfolio of variant conditions — thus the tendency to play whack-a-mole.

Decision-making

Because decision-making is so fundamental to effective risk management, it is important to ensure that decisions are made by the appropriate personnel. Of course, as emphasized throughout this maturity model, ensuring that decision-makers are provided accurate and meaningful risk information is equally important.

Many organizations have begun to pay attention to who's making risk decisions by formally defining the scope of decision-making authority. The next step in maturity is to periodically review decisions (policy exceptions, risk acceptances, etc.) to ensure that those authorities aren't exceeded.

Few organizations, however, actively and purposefully examine the quality of risk information being provided to decision-makers. Too often, it is simply assumed that when somebody proclaims something to be high/medium/low risk, they got it right. Unfortunately, that assumption is rarely well-founded. Organizations that take risk measurement and management seriously will periodically have quality reviews performed on risk analyses. These reviews should be performed by an independent party, which can be internal if they have such expertise in-house.

Analysis

If visibility provides data, then analysis provides intelligence. Regardless of how complete an organization's visibility is (or isn't), organizations must be able to evaluate their available data in order to accurately determine what the risk conditions mean from a business/organization perspective.

Arriving at this risk-related intelligence from data invariably requires two elements — models used to evaluate the data, and the analysts who apply the data to the models. Unfortunately, few organizations analyze their risk-related data with any rigor. The wet finger in the air predominates. In this mode, both the models being used and the personnel doing the "analysis" are highly suspect. Another common approach is to rely on "risk ratings" provided by some security technologies (e.g., vulnerability scanners). Today however, these risk ratings are rarely based on strong models. The resulting inaccurate risk ratings tend to overwhelm organizations with "high risk" issues that aren't truly high risk. This makes it exceptionally difficult to identify and focus on the things that matter most from a risk perspective.

Models

Every risk measurement involves the use of a model. At the low end of the maturity continuum, the model being used is an individual's uncalibrated mental model based on whatever experience, bias, and education they bring to the table. At the other end of the continuum, it's a formal model that has been validated through independent evaluation and/or empirical testing. In between, of course, are things like the calibrated mental models of trained risk analysts, and formal models that have undergone less rigorous validation.

Analyst skills

Extremely few organizations have personnel who are dedicated to risk analysis. In most cases, organizations simply rely on personnel who are experts in one or more aspects of cyber risk (e.g., auditors, network security architects, application security professionals, threat intelligence professionals, etc.) to rate the significance of risk-related conditions. Although these professionals may be brilliant in their area of expertise, they rarely have (or at least rarely apply) the skills necessary to reliably analyze and measure risk.

Being adept at risk analysis requires a specific set of characteristics, which includes:

- Strong critical thinking skills
- The ability to view complex problems from multiple perspectives

- Being comfortable with numbers
- Understanding of, and ability to apply, basic probability principles
- Being comfortable with uncertainty — i.e., does not view the world as black and white, or freeze up when data are incomplete or imperfect
- Trained in making calibrated estimates
- Trained in the application of models — e.g., scoping analyses

Absent these characteristics, the odds of a person applying data and models to accurately measure risk are considerably diminished.

It's important to recognize, too, that organizations cannot afford to analyze every possible issue to the nth degree. There just aren't enough resources. Therefore, analysts must be able to prioritize their analysis efforts through a simplified analysis (triage) process.

Reporting

If analysis provides intelligence, then reporting provides decision-makers with understanding. In other words, analysis results can be profoundly confusing or meaningless unless they are conveyed in a way that decision-makers can understand and act on.

Typical cyber risk reporting today includes one or both of the following:

- Technical metrics related to things like attack volumes, patch levels, and awareness training.
- Heat maps of the organization's top "risks"

In the first case, technical metrics are rarely meaningful at an organization leadership level. Usually, the decision-makers are left to imagine for themselves what the information means from a business risk perspective, or they rely on the security/audit professional to interpret for them. Neither usually results in an accurate understanding of risk.

In the second case, heat maps usually fail to provide business relevance. What does "high risk" mean from a business perspective? Furthermore, which of the high-risk issues is most significant and by how much? Worse yet, many of the "risks" in these reports aren't risks. A previous white paper on Risk Clarification provides additional information regarding this problem.

In either case, decision-makers are challenged to appropriately balance cyber risk concerns with the other business imperatives (e.g., opportunities, operational concerns, and other forms of risk). Given the fact that organizations must somehow balance their limited resources across all of these imperatives, any deficiency in making good choices diminishes an organization's odds of overall success.

Reliable execution

Well-informed decisions are simply half of the risk management equation. The other half is reliable execution in alignment with those decisions. In this section, we'll discuss the three elements that contribute to reliable execution.

Awareness

If personnel responsible for execution aren't clearly aware of what's expected of them, then the odds of reliable execution are low. Too much is left to chance. This is just as true for the sales person who's not supposed to send sensitive information via e-mail, as it is for the mid-level manager who is supposed to notify Identity and Access Management when personnel change roles. For that matter, it also applies to software developers who are expected to write secure code, and decision-makers who need to know the level of risk they are authorized to accept.

Unfortunately, although many organizations put significant effort toward educating personnel in the mom-and-apple-pie basics of good password selection, anti-phishing, and clean desktops, few organizations put significant effort into targeted education for developers, system admins, line managers, executives, etc. The resulting superficial awareness levels increase the odds of misaligned execution of responsibilities specific to those roles.

Capability

Even when personnel are aware of what's expected of them, they still must have the skills and resources to perform reliably. This very often boils down to whether personnel have the necessary training and tools to reliably fulfill their responsibilities. Training, in particular, can be problematic because when budget belts get tightened, training is often one of the first things on the chopping block. Another common sacrifice, of course, is simple bandwidth — i.e., reduced staffing levels can strongly inhibit the ability for personnel to keep up with their responsibilities. When this happens, uneven incentives can become an even bigger factor, as discussed below.

Motivation

Basic economics has taught us that incentives drive behavior. For example, personnel throughout many organizations are often strongly and formally incentivized to exceed revenue or mission-related goals, and not to exceed budget limits or project deadlines. Rarely is anyone outside of the cyber risk organization formally incentivized to hit cyber risk management objectives. As a result, when push comes to shove, risk management priorities tend to get the short end of the stick. For cyber risk concerns to compete on a level playing field, they must be equally and formally emphasized by leadership.

Ultimately, reliable execution is predicated on well-informed decision-makers. In other words, decision-makers can only make appropriate adjustments to education, capabilities, or motivation if they have accurate information about execution deficiencies and the risk implications of those deficiencies.

An example set of maturity measurements

Over the past year, a version of this maturity model has been used to evaluate numerous organizations. In many cases, these organizations had scored well on more traditional frameworks like the PCI DSS, NIST CSF, etc. To-date, however, no organization has scored well against this maturity model. The implication is that although organizations spend significant time and money on cyber-related controls, relatively little attention is being paid to ensuring that well-informed decision-making or reliable execution takes place.

To help keep this maturity model simple, each element is evaluated using a three-level scale. An attempt has been made to keep the scale descriptions as unambiguous and objective as possible. Regardless, some ambiguity and subjectivity inevitably remains. Like all models, this one is continually evolving based on research and testing.

Risk terminology

Strong

A standard set of risk-related terms has been formally defined or adopted. Personnel within the risk management organization (including cyber risk, audit, privacy, compliance, technology, operational risk, etc.) understand and consistently apply these terms. Inconsistent usage is corrected.

Partial

A standard set of risk-related terms has been formally defined or adopted, but usage is inconsistent.

Weak

No standard set of risk-related terms has been defined or adopted. If you ask six people in the risk management organization to define foundational risk-related terms or provide examples of what those terms represent, you will likely receive different answers.

Asset Visibility

Strong

An inventory of systems, applications, and significant information repositories exists and is kept up-to-date through well-defined and consistently practiced procedures. An audit of the inventory would be unlikely to find that more than 5% of the entries are inaccurate. Also, the value/liability characteristics of assets (e.g., classification) is included in the inventory.

Partial

An inventory of systems, applications, and significant information repositories exists but is not consistently maintained. Processes for maintaining the inventory are immature or are exercised unreliably. Audits of the inventory regularly find more than 5% of the entries are inaccurate.

Weak

An inventory of systems, applications, and significant information repositories does not exist or is severely out of date (i.e., cannot be relied on to support decision-making). Processes for maintaining the inventory either do not exist or are not practiced.

Threat Visibility

Strong

Threat intelligence is a specialization within the information security group (or has been outsourced) and is capable of providing organization-specific information regarding changes in the threat landscape (e.g., increases in the frequency or sophistication of attacks experienced by the organization). Threat data for key assets and points of attack are closely monitored.

Partial

Threat intelligence is received from internal resources and/or external sources (e.g., ISAC organizations) that provide information regarding changes and trends in the general threat landscape (e.g., the existence of a new zero-day exploit) as well as the organization's industry.

Weak

Threat intelligence is acquired in an informal or ad hoc manner (e.g., blogs, mailing lists, etc.) and is highly generalized in nature (i.e., the data is not specific to the organization's industry or the organization itself).

Controls Visibility

Strong

The frequency of controls testing (e.g., authentication, access privilege, configuration, and patch conditions, etc.,) is driven by the value/liability characteristics of the assets, the level of threat they face, and the anticipated degree of change surrounding those assets. In other words, controls testing is more frequent for assets that are of higher value, face a more active threat landscape, and that undergo more frequent changes.

Partial

Authentication, access privilege, configuration and patch conditions, etc., are tested on a regular basis but the testing regimen is not risk-based. As a result, some key systems, applications, or points of attack may not get tested at all or testing occurs infrequently.

Weak

Authentication, access privilege, configuration, and patch conditions, etc., are infrequently tested and not well known.

Execution Visibility

Strong

Root cause analysis of non-compliant conditions is performed at least 75% of the time when non-compliant conditions are discovered. The population of root-cause analyses are evaluated as a portfolio to discover systemic sources of non-compliance.

Partial

Root cause analysis is periodically performed (at least half the time) when non-compliant conditions are discovered, but no attempt is made to perform a portfolio review of these analyses in an attempt to discover systemic problems within the organization.

Weak

Root cause analysis is not performed (or is performed less than half the time) when non-compliant conditions are discovered.

Decision-making Visibility

Strong

At least once per year the organization performs both of the following: 1) reviews risk management decisions to ensure that they are being made at the appropriate level of leadership, and 2) has an independent review performed of risk ratings/values to validate that the risk information being provided to decision-makers is accurate.

Partial

At least once per year the organization performs one of the following: 1) reviews decisions (e.g., policy exception requests, policy/standards development, etc.) to ensure they are being made by the appropriate personnel, or 2) has an independent review performed of risk ratings/values that were used to validate that the risk information being provided to decision-makers is accurate.

Weak

The organization does not review risk management decision-making to ensure that decisions are being made by the appropriate personnel or that risk measurements were accurate.

Models

Strong

Risk analyses consistently leverage a well-defined and publicly vetted analytic framework (i.e., is not checklist-based). An example would be the OpenFAIR model.

Partial

Risk analyses rely on models that have been developed internally or by a third party, and that have not undergone independent validation.

Weak

Analysis relies primarily on the intuition (mental models) of subject matter experts. Little or no documentation or validation of the underlying assumptions takes place.

Analyst Skills

Strong

The organization has (or contracts to) personnel who are dedicated to performing risk analysis. Analysts have expertise in quantitative risk measurement concepts and principles, and have been specifically trained in the process of scoping scenarios and making calibrated estimates.

Partial

Analysts are not dedicated specifically to performing risk analysis. They have experience in performing qualitative information security risk analyses, but may have limited expertise in formal analysis methods, probability principles, etc.

Weak

Analysts are experienced in information security and/or technology but are inexperienced in formal risk analysis methods.

Reporting

Strong

Risk reporting includes quantitative statements of risk so that decision-makers can effectively compare and prioritize information security concerns against other organization concerns (e.g., operational needs, growth opportunities, and other forms of risk).

Partial

Risk reporting is worded for the intended audience but is primarily qualitative in nature.

Weak

Risk reporting to operational and executive management contains a significant amount of technical information and jargon.

Awareness

Strong

The organization has documented and published policies, standards and processes and these documents are kept up-to-date. Personnel are required to understand the specific risk management expectations for their job responsibilities (e.g., developers understand secure software standards, system and network administrators understand configuration, change management, and architecture standards, etc.) and their understanding of these expectations is evaluated once per year.

Partial

The organization has documented and published policies, standards and processes and these documents are mostly kept up-to-date. Personnel are required to read and acknowledge their understanding of the organization's general risk management expectations.

Weak

The organization has little or no documented and published policies, standards, and processes, or these documents are out-of-date. There are no active processes in place to make personnel aware of these expectations. Most personnel have little or no understanding of the organization's risk management expectations.

Capability

Strong

Updated training in relevant risk management areas of expertise is required on an annual basis to help ensure that personnel keep abreast of changes in the risk landscape, technology, and/or processes. Funding for this effort is not subject to budget cuts.

Partial

Updated risk-related training is typically provided but not required. Funding for training is subject to budget cuts.

Weak

Updated training is not provided or is inconsistently funded.

Motivation

Strong

Key cyber risk objectives are formally defined within the performance expectations and compensation/bonus plans for senior business leadership. Failing to meet cyber risk objectives consistently results in the same (or more severe) consequences as failing to meet revenue goals, exceeding deadlines, exceeding budget limits, etc.

Partial

Cyber risk objectives are included in the performance expectations/reviews for key personnel with risk management responsibilities (e.g., system admins, software developers, etc.). Failing to meet cyber risk objectives can result in the same (or more severe) consequences as failing to meet deadlines, exceeding budget limits, etc.

Weak

Failing to meet cyber risk expectations/objectives rarely results in meaningful consequences.

A note regarding this model

RiskLens has developed a Cyber Risk Maturity software application that uses a version of this model. The application leverages a Bayesian network as the underlying analytic engine to capture and reflect relationships between the model elements.