CHAPTER

Information Security Metrics

13

Would you believe us if we told you there was one metric, and only one, that would tell you everything you needed to know about an organization's information security risk posture? No, probably not, and you'd be right. That said, the number of metrics required to gain a meaningful understanding of an organization's risk posture is not hundreds, or even dozens. Not, at least, if you understand the key elements that drive risk into an organization.

A comprehensive coverage of metrics isn't possible in a single chapter. Therefore, our focus here is on understanding where and how metrics fit into an information security risk management program and how to leverage them effectively. To provide this understanding, we'll share a framework you can use for identifying meaningful metrics and figuring out their value proposition. Of course, we'll provide examples, both in this chapter and the next one. Our examples will not be exhaustive, though. If you're looking for exhaustive examples, there are a lot of books on the market dedicated to information security metrics.

Speaking of other books, if you're looking for a very good book dedicated to the subject of information security metrics, we really like *IT Security Metrics* by Lance Hayden. Hayden goes into significant detail on the nature of data, statistics, and analysis. For the data geeks in the crowd, we also really like another book entitled *Data-Driven Security: Analysis, Visualization, and Dashboards* by Jay Jacobs and Bob Rudis. We would like to think that the concepts and frameworks presented in this book, and especially this chapter, will allow you to better leverage the wealth of information in those books.

TALKING ABOUT RISK

The word "data" can take a plural or singular form (e.g., "data are" or "data is"). Scientists and other quants often prefer the plural form, while much of the rest of humanity seems to prefer the singular form. Fortunately, the word's meaning does not change based on the form that is used, nor is there confusion about what it means. For those reasons, we didn't sweat it. We hope you wouldn't either, at least while reading this book.

CURRENT STATE OF AFFAIRS

The use of metrics promises to take our profession from art to science (or at least to something less superficial and more science-like). In order to realize that promise, however, our profession has to solve a few fundamental problems first—problems

Measuring and Managing Information Risk. http://dx.doi.org/10.1016/B978-0-12-420231-3.00013-0

we have beaten a drum about throughout this book. For example, without consistent and logical nomenclature, it becomes wickedly hard to normalize data or communicate effectively. After all, if one person's "threat" is another person's "risk" is another person's "vulnerability," it is extremely difficult to find common ground. How do you know what data you need in the first place, and how do you apply data to derive meaningful results if your "models" look anything like "Threat x Vulnerability/Controls," or are simply checklists? Finally, the only way your metrics become meaningful is if they support explicitly defined objectives that matter. In this chapter, we continue the process of tying together what has been covered in the earlier chapters—nomenclature, models, and objectives—so that you can leverage metrics more effectively.

TALKING ABOUT RISK

We have heard the statement on more than one occasion that an important criterion for a "good" metric is that the data should be easy to acquire. Yes, it's great when data acquisition is easy. but if you rely on that to drive which metrics you use, you may miss out on really important information. All we're saying is don't just rely on the easy stuff. Understand the decision you are trying to support and get the best information you can, given your time and resources.

TALKING ABOUT RISK

It's been our experience that information security organizations can often get away with having relatively useless (or worse, misleading) metrics. On numerous occasions, we have seen auditors, regulators, executives, and third-party assessors apparently attribute program maturity and effectiveness to a bunch of colorful charts and graphs, even when the metrics are either misleading or go entirely unused in decision-making.

METRIC VALUE PROPOSITION

Remember what we said in the Risk Management chapter; that risk management boils down to a series of decisions and the execution of those decisions? Well, this entire chapter could perhaps be entitled "Decision Support" because the only reason for generating metrics is to inform decisions. In fact, *if you're publishing metrics that aren't being actively used in decision-making, then you are wasting time and resources*. Because of this, we're going to come at metrics with a clear eye toward their role in decision support. Behind every decision there are one or more goals that an organization is driving toward. Before we go on, ask yourself what overarching goal might form the foundation for decisions within a risk management program. We'll answer the question shortly, but here is a hint—we discussed it in the Risk Management chapter.

Within the metrics world, you may have heard people talk about the "Goal, Question, Metric" (GQM) method for developing good metrics. We really like this approach because it helps people focus on and understand a metric's value proposition[1]. For example, the GQM approach might go about defining a metric in the following way:

- Goal: Reduce the number of network shares containing sensitive information
- Question: How much sensitive information resides on network shares?
- Metric: Volume of sensitive information on network shares

This is a clear and concise way to define that kind of metric. However, you have to be a bit careful to not put the cart before the horse. The above example suggests that a decision had already been made regarding a different question. Perhaps that different question was, "Do we need to reduce the volume of sensitive information on our network shares?" (Apparently, the answer was "yes"). There may have been a question before that; something like, "Do we have significant concentrations of risk associated with sensitive information?" (Again, apparently the answer was "yes"). Absent the context of those questions and their subsequent decisions, chasing a metric like the volume of sensitive information on network shares might not be a good use of time even given a great metric definition method like GQM. We have to define the big picture—those "macro goals"—first.

In keeping with our decision-based focus, we would like to make a fairly subtle but important observation about the question component of GQM. In the above example related to network shares, the original question was phrased as "how much," yet the implied questions that might have come before were phrased differently. The "Do we need…?" and "Do we have…?" phrasing is more explicitly aligned with decision-making because, depending on the answers, different actions may be required. The question of "how much" doesn't explicitly relate to a decision or goal. Implicitly, perhaps, but it's important that we understand the decision context for the metric as explicitly as possible.

Before we dive into the section on how to leverage GQM to make metrics meaningful, there is one more thing to point out—comparison. Specifically, metrics are fundamentally a means of making comparisons between, for example:

- Current conditions and desired future conditions
- Risk scenarios (prioritization)
- Mitigation options (selection)
- Past conditions and current conditions (efficacy of past decisions and actions)

You may have noticed that this also aligns with the risk management stack—meaningful measurements enable effective comparisons, which enable well-informed decisions. We love it when things come together like this.

BEGINNING WITH THE END IN MIND

So, did you come up with any ideas regarding our question about an overarching GQM-type goal for metrics? As you'll recall from the Risk Management chapter, our

[1] The IT Security Metrics book by Lance Hayden does a great job of discussing GQM.

definition for risk management includes the phrase, "*…cost-effectively achieve and maintain an acceptable level of loss exposure*." That sounds suspiciously like a goal to us, so you get a diamond encrusted platinum star if that's what you came up with. With that goal as our starting point, let's continue to break this down and apply the GQM approach for our metrics. We can begin by breaking our overarching goal into four subgoals:

- Being cost-effective
- Achieving alignment with the organization's risk appetite
- Maintaining alignment with the organization's risk appetite
- Defining the organization's risk appetite (that "acceptable level of loss exposure")

The next step is to break these down into more granular subgoals.

BREAKING IT DOWN

In this section, we'll begin to break down our overarching goal into layers of subgoals, questions, and metrics. We'll wait to discuss these subgoals until a little further on though, because some of the discussion will be lengthy. For now, let's just cover the outline.

In the mind-map shown in Figure 13.1, we have broken out our four main subgoals into another layer of granularity. Once you have a handle on this layer of abstraction, we think you will find it is pretty easy to figure out additional layers of subgoals, questions for these goals, and then the metrics that inform those questions and their associated decisions.

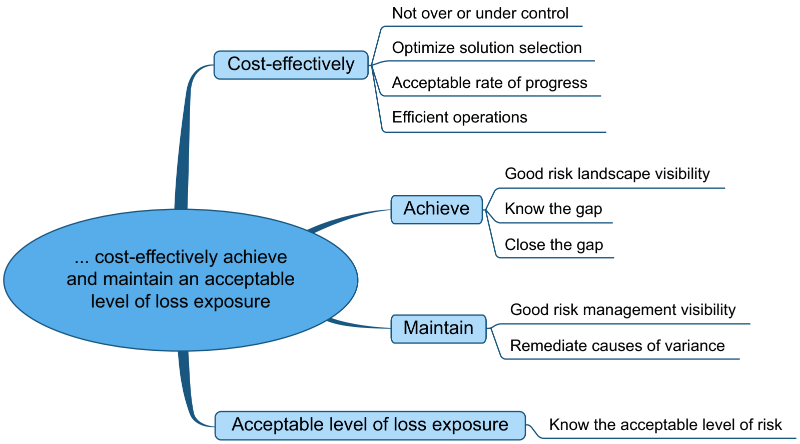
... cost-effectively achieve and maintain an acceptable level of loss exposure
Cost-effectively
Not over or under control
Optimize solution selection
Acceptable rate of progress
Efficient operations
Achieve
Good risk landscape visibility
Know the gap
Close the gap
Maintain
Good risk management visibility
Remediate causes of variance
Acceptable level of loss exposure
Know the acceptable level of risk

FIGURE 13.1

Mindmap of example risk management metric goals.

Before we move on, there are a few things we need to point out about the framework above:

- This framework may not be the only way to decompose the risk management goal. You may find that a different set of subgoals, questions, and metrics works better for you, or you might find that this framework provides most of what you need, only requiring a handful of your own tweaks.
- Efforts for achieving and maintaining will likely (should likely) run in parallel. If you aren't tackling the root causes for variance and unacceptable loss exposure (primarily a part of the "maintain" function) even as you are mitigating current exposure, then your progress in achieving an acceptable level of risk is likely to be much slower. This is because even as you fix things, bad risk management practices will be introducing risk elsewhere.
- Even if your organization achieves its desired level of risk, the dynamic nature of the risk landscape is undoubtedly going to throw occasional curveballs that take the organization out of that comfort zone. Also, keep in mind that in some cases, management's comfort zone may shift. Either way, you should view this as a never-ending process. However, that has always been a mantra of information security and risk management, right?

The bottom line is that the value of any metric should be defined within the context of a goal. When it comes to information security, that goal is to manage risk cost-effectively over time through better-informed decisions. Starting with that as the focus, the rest is easy; or at least easier.

MORE DETAILS

Okay, let's dig in and start to flesh this framework out a bit. As we do, you might very well identify questions and metrics (maybe even subgoals) that you would like to add. If so, go for it. This table (Table 13.1) is not intended to provide the final word on the topic or be entirely comprehensive. It's intended to demonstrate the approach and provide a good starting point. You will note, too, that some metrics will support more than one question or goal.

After this section, we'll discuss each of the subgoals in greater detail to help ensure clarity and because some of them are particularly challenging in nature (e.g., determining an organization's acceptable level of loss exposure).

One thing we hope you've noticed is that most of this is dependent on the ability to measure risk. We'll discuss this in greater detail further on, but we didn't want to miss the opportunity to drive home the point that effective risk management, particularly cost-effective risk management, is naturally dependent on risk measurement practices. If risk measurement isn't meaningful or is done poorly, then the odds of effectively managing risk go way down.

You may note that we didn't explicitly include examples related to reducing controls if, for example, an organization found it was overcontrolling loss exposure in some part of its risk landscape. Keep in mind though that the "close the gap" goal

Table 13.1 Example Goals, Subgoals, Questions, and Metrics

Goals	Subgoals	Questions	Metrics
Cost-effective	Not over or under control	Are we aligned with risk appetite?	Acceptable level of risk Current risk level
	Optimize solution selection	What is the most cost-effective solution?	Solution costs Solution benefits
	Acceptable rate of progress	Are we progressing toward objectives at the proper rate?	Milestones Current risk condition Previous risk condition Elapsed time Forecast risk condition
	Efficient operations	Are we focused on the most important things?	Areas of risk concentration Key control deficiencies
		Is the full cost-benefit of our resources being realized?	Resource utilization Resource cost
Achieve	Good risk landscape visibility	Do we have good visibility?	Threat intelligence Asset management Control conditions Impact factors
	Know the gap	How far away from alignment are we?	Acceptable level of risk Current risk level
	Close the gap	Where does risk exist?	Risk assessments Self-identified points of exposure Loss events
		What control deficiencies exist?	Risk assessments Self-identified deficiencies Loss events
Maintain	Good risk management visibility	Do we have good risk management visibility?	Asset visibility Threat visibility Controls visibility Impact factor visibility Decision visibility Execution visibility
	Remediate causes of variance	Which root causes are driving variance into the environment?	Variance data Root cause analysis results
Acceptable level of loss exposure	Know the acceptable level of loss exposure	What is the acceptable level of loss exposure?	Current risk level The organization's loss capacity Management's tolerance for loss

does not specifically say, "reduce risk." In some cases, closing the gap might mean, "increase risk." This is something to keep in mind because we have in our careers found numerous instances where controls (e.g., a lot of the common SOX controls, as an example) could be relaxed to significantly reduce the business burden and not result in a material increase in risk. As you might imagine, when you find these opportunities, management loves you for it.

In this next section, we will begin to examine our subgoals and metrics in more detail. For reasons that will become obvious, we'll cover them in a different order than we have presented them above.

TALKING ABOUT RISK

Since we brought it up… I (Jack Jones) have in the past been something of a nuisance to external and internal auditors on the subject of SOX controls. In one organization, I challenged them to describe the relevance of many of the common IT SOX controls outside of a science fiction loss scenario. For most of the controls in question, they couldn't. As a result, we worked together to review the SOX controls, ended up eliminating some, and demoting a significant percentage of the others to general IT controls status. This resulted in more accurate SOX-related risk reports for management, fewer wasted resources dealing with burdensome SOX processes on controls that didn't warrant it, and a very happy set of executives. Cost-effective risk management in action!

ACCEPTABLE LEVEL OF LOSS EXPOSURE

We need to tackle this goal first because the other goals are inherently dependent on this one. What we're talking about here is risk appetite, or tolerance, depending on who you're talking to (more on this later). Regardless of what you call it, what it boils down to is the level of loss exposure executive management is comfortable with. This is an aspect of risk management that information security organizations rarely tackle effectively, if at all.

Perhaps the most common approach we've seen (and we have approached it this way in the past) is to assume that an organization's information security policies are a tangible representation of management's risk appetite. Based on this assumption, measurements focus on variance from policies and standards. This approach is relatively straightforward and conceptually correct, but it is also inherently implicit and has a couple of fundamental flaws, at least as it is most commonly done:

- It is relatively unusual to encounter organizations whose policies and standards have actually been reviewed in a meaningful manner with executive management. More commonly, the information security team defined (or downloaded, or had a consultant define) the policies and standards and had management "sign-off" on them. Unfortunately, if the risk management value and cost implications of policies and standards aren't discussed with management, they won't know what they are signing-off on. As a result, the policies and standards may not actually set the stage for a level of risk, and cost, that management is okay with, which increases the probability of getting wrapped up in the Groundhog Day phenomena we mentioned earlier.

- Risk status reports based solely on compliance levels still aren't inherently meaningful to management. If the organization isn't complying effectively with one requirement or another, so what? How much should management care?

The bottom line is that if your organization operates on the assumption that policies and standards represent the organization's risk appetite then it is critical that you do as much as possible to make sure the assumption is accurate. As we described in the previous chapter, getting face time with executives to discuss policy risk and cost implications is a critical first step. Another critical requirement is to monitor policy exceptions and enforcement practices. If you see a lot of exceptions to a policy, or inconsistent enforcement of violations, it is a good sign that the policy doesn't match how management wants to operate (either that, or management has not been given the information they need to understand the risk implications). Finally, it is also important to ensure that when reports regarding noncompliance are taken to management, the "so what" question is answered through meaningful risk analysis. This enables more explicit risk management.

A more effective approach to defining risk appetite in many organizations is to define it in monetary terms. That's how it's done in the financially focused risk disciplines like credit risk and investment risk, and it is inherently more meaningful to executive management. Unfortunately, information security loss exposure has traditionally been expressed as a set of high, medium, and low conditions (usually related to control deficiencies versus loss event scenarios). It's a significant leap of faith to believe that high, medium, and low terms are really meaningful to an organization's executives. Besides, how many "highs" are too many? How many "mediums" equal a "high?" What do we mean by "high risk", anyway? What are the loss event scenarios to which those ratings are referring? This being the case, information security has been badly hamstrung in defining the most foundational element of its overall goal! Without that definition, the ability to leverage GQM in a truly meaningful manner is likewise hamstrung.

The information security discipline faces another challenge though, beyond measuring risk in meaningful terms. In credit risk, an organization is (or should be) able to keep reasonably good track of the credit it has given and taken, and define a threshold for the related loss exposure that it doesn't want to exceed. Unfortunately, in information security it is arguably impossible to have complete visibility into all the nooks and crannies where things can and do go wrong. This visibility problem makes it much more challenging to represent the complete picture of loss exposure an organization faces, and thus precisely define its acceptable loss exposure threshold. That said, the problem is not insurmountable as long as several points are kept in mind:

- Visibility needs to be an explicit element of your risk management metrics so that areas of poor visibility are recognized and improved on over time.
- The overall confidence in aggregate loss exposure measurements should be tempered by the organization's risk landscape visibility.
- We should not expect high degrees of precision in aggregate loss exposure measurements. Risk measurement and metrics is about getting better, more

useful information than the typical high, medium, and low ratings (which, keep in mind, suffer from the same visibility problem).

We'll spend a lot more time on this question of visibility later in the chapter.

TALKING ABOUT RISK

The terms "risk appetite" and its close cousin "risk tolerance" are often poorly understood, very rarely used to good effect, and commonly used interchangeably. Similar to the word "risk," you will sometimes get as many different definitions for these terms as people you ask. Potentially useful definitions we have seen include:

- Risk appetite: A target level of loss exposure that the organization views as acceptable, given business objectives and resources
- Risk tolerance: The degree of variance from the organization's risk appetite that the organization is willing to tolerate

Given these definitions, a simple analogy for appetite and tolerance would be speed on a highway. The department of transportation or other government entity sets a speed limit. This could be roughly thought of as analogous to risk appetite and reflects the decision-makers beliefs regarding an appropriate balance between traffic flow, highway and environmental wear-and-tear, and public safety (among other things). The people using the highway will usually travel at speeds greater or lesser than the speed limit as opposed to exactly at the speed limit, and the point at which law enforcement actually begins ticketing violators could be viewed as analogous to risk tolerance. Given normal weather and other conditions, it is extremely rare to see law enforcement enforce the speed exactly at the limit. Consequently, while risk appetite can be thought of as a line drawn in the sand that helps to set expectations, risk tolerance can be thought of as the variance from appetite that drives day-to-day decisions to operate differently in some manner. Note the operative word here—decisions.

Defining an acceptable level of loss exposure

If we want to define an acceptable level of loss exposure for an organization, how do we go about doing that? Our GQM framework tells us that we need three pieces of information: the organization's current level of risk, the organization's capacity for loss, and management's tolerance for loss. Of these three elements, information security can only provide one—the organization's current level of risk. The other two elements are pieces of the puzzle that executive management brings to the table and may not, in fact, state explicitly.

TALKING ABOUT RISK

Why do we need the organization's current level of risk when establishing risk appetite? Actually, we shouldn't need it. All we should need is the organization's capacity for loss and the executive management's personal views on risk. That said, there is a decent chance today that if you ask executives to give you a specific information security risk appetite value, they will say something like "less" and shrug—ironic, eh? Many of them simply haven't considered the possibility that information security risk can be quantified and aggregated because all they have ever seen are red, yellow, and green heatmaps. They may, of course, decide to reference a quantitative risk appetite definition that the organization uses for one of its financial risk domains, which is not a bad thing. Or, perhaps, to paraphrase Supreme Court Justice Potter Stewart's answer when he was asked why a particular motion picture did not qualify as obscene, an executive may respond with, "*I'll know it when I see it*."

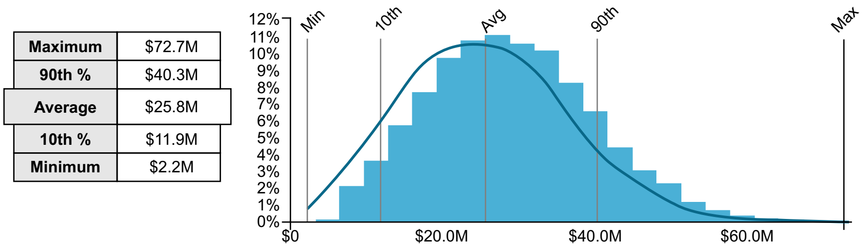

Maximum	$72.7M
90th %	$40.3M
Average	$25.8M
10th %	$11.9M
Minimum	$2.2M

FIGURE 13.2

Example aggregate risk distribution.

If the organization's current level of loss exposure is an important piece of information that's needed to set risk appetite, how do we go about measuring it? There are a couple of ways we have seen that work well.

One option is for an organization to perform a high-level triage of the results of its most recent security review or audit. Assuming the examination was comprehensive and identified the control conditions that were significantly different from what they were supposed to be (i.e., were variances), a quick, risk-based triage of those findings can be performed to identify which of those conditions and related scenarios appear to represent the greatest exposure. Having applied that filter, full-fledged factor analysis of information risk (FAIR) analyses can then performed on these scenarios. Now, to get the big picture you can't simply sum up the loss exposure from each of these analyses, as that would represent the loss exposure if all of moons aligned simultaneously (kind of an asteroid-strike view of the world). You have to apply Bernoulli and/or Poisson functions by using a spreadsheet, statistical tool, or (simpler still) a commercial FAIR application (see Figure 13.2).

This will give you a distribution of joint probabilities that provides a truer view of the aggregate exposure from those scenarios. It's only a partial view of the overall loss exposure of the organization, but if the security review/audit and analyses were done well, then the results should provide a good approximation of where the organization stands with loss exposure. It is certainly a great place to start. If you prefer a simpler approach (but one that is a couple of steps lower in terms of effectiveness and explicitness), you can plot the FAIR results on a heat-map.

TALKING ABOUT RISK

We have found the question of incomplete visibility and its effect on analysis results to be a great talking point with management. Specifically, if visibility into parts of the risk landscape is weak, and if that is believed to impact the quality of a risk analysis and leadership's ability to make better-informed decisions, then management is more likely to support visibility improvement efforts.

Senior executives who have had to deal with setting thresholds in other areas of risk should be able to look at this information and come to a reasonable conclusion about whether that amount of loss exposure is, in their eyes, acceptable for the organization. Is the analysis precise? Of course not. As we've discussed elsewhere, high levels of precision aren't the goal. The goal is more useful information than they've had before. We find it hard to believe that an executive would prefer a list containing five "high risk" issues, and 24 "medium risk" issues over the quantitative data above. Furthermore, with the approach above, you're also able to identify which scenarios, assets, and control deficiencies are contributing most to the overall level of risk, which can always be plotted on a heat-map to help with prioritization.

Another approach is to use something like CXOWARE's RiskCalibrator application, which significantly simplifies the process. Think of it as a TurboTax-like approach to performing an enterprise information security risk analysis that automatically provides the aggregate loss-exposure picture, as well as identifying where risk is concentrated, which threat communities are most problematic, and which control deficiencies are driving the most risk into the organization (and thus where the best bang-for-the-buck is from a remediation perspective). It also performs sensitivity analysis so that the organization knows where to focus on improving metrics quality, and where the organization is sensitive to changes in the risk landscape in either good or bad ways.

Imagine for a moment that you are incredibly fortunate (from our point of view) and you work for an organization that has established a quantitative statement of information security risk appetite. Congratulations! Now what? Even if you performed an analysis like we described above, you have at least one more important question to answer. Which part of the analysis results do you compare against the defined threshold? You see, the outcome of an aggregate risk analysis is not a discrete number (e.g., $X of loss exposure). No, the outcome is a distribution, like the one shown in Figure 13.3, of potential levels of loss exposure reflecting the imprecision and inevitable uncertainty in the analysis. So which number from within this distribution do you use for comparison against the defined risk appetite? The average? The maximum? One of the percentiles? Or do you use the single event worst-case? The unsatisfying answer to this question is, "It depends."

Suppose your organization defined its information security risk appetite as $30 million of annualized exposure. If we compare that value against the aggregate

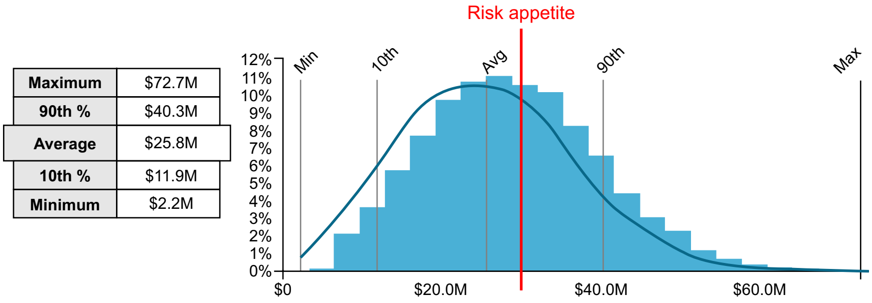

Maximum	$72.7M
90th %	$40.3M
Average	$25.8M
10th %	$11.9M
Minimum	$2.2M

FIGURE 13.3

Risk appetite.

results shown in the chart above, it falls in between the average and the 90th percentile. So, does the organization's aggregate loss exposure fall within its defined appetite or not? Well, the answer is "yes" if the organization has chosen to reference against the average, and "no" if the organization has chosen to reference against the 90th percentile. The point is, an organization needs to establish an agreed upon expectation regarding which point of a loss exposure distribution will be used for such comparisons. Where on the distribution an organization settles also provides a hint regarding its risk tendencies. An organization that chooses to reference against the 90th percentile of the distribution can probably be thought of as more risk adverse than one that references against the average.

Another possibility is that management won't set a specific risk appetite value, but instead will look at the organization's current loss exposure and say something like, "What will it take to cut that in half in the next 24 months?" That's fine, too. It's specific, which enables explicit risk management to take place.

TALKING ABOUT RISK

If you are thinking about talking with your organization's senior executives about defining a quantitative information security risk appetite, you should be aware that they are very likely to ask the question, "Who else is already doing this?" and "Is there a benchmark we can compare ourselves against?" The unfortunate fact (at the time of this writing) is that the answers to these questions are, "almost nobody" and "no." FAIR has been plowing new ground regarding measuring information security risk for a dozen years now, and has gained recognition and use in tactical decision-making. Using quantitative risk analysis in the strategic management of information security risk is the next logical progression. We believe it's simply a matter of time before this gains widespread recognition and adoption.

An important thing to keep in mind regarding "an acceptable level of loss exposure" is that it is highly situational. Although an organization may in normal circumstances enforce relatively close alignment to its defined risk appetite, opportunities or constraints may create circumstances where organization leadership knowingly and willingly make decisions regarding risk that takes the organization outside of the defined threshold. This may take the form of additional risk-taking or reduced risk-taking. This is where the notion of risk tolerance could come into play, at least if you subscribe to the definition we mentioned earlier. An organization might define thresholds above and below its normal risk appetite, beyond which a careful examination and exceptional approvals might need to occur regarding risk decisions, policies, etc.

TALKING ABOUT RISK

There is something that strikes us as odd about the belief held by some information security professionals that executives have an almost unbounded appetite for information security risk. It's odd because studies show that when people are given a choice between potentially gaining something versus potentially losing something of similar value, they will choose to avoid the loss. As a species, humans are simply wired to be risk averse. In our evolutionary past, although the need to acquire food, water, mates, etc. was important to survival, we could (up to a point) continue to survive if we passed up on those opportunities. If, however, we failed to avoid a significant threat, then we died. Game over, plain and simple. As a result, we are wired to more highly prioritize loss avoidance. In order for people to choose the “gain” option, they have to believe that the potential for gain is meaningfully greater than the potential for loss. Consequently, if an information security professional’s experience is that executives routinely disregard warnings of “high risk,” then that suggests one or more of the following:

- Either the information security professional did not effectively communicate the loss exposure,
- The executives do not view the information as credible, or
- The risk information they are putting in front of the executives truly is not relevant when compared to the opportunities or competing risk issues the executives have to deal with

Our experience has been that when we present risk in terms that are meaningful, that we can defend, and that are truly significant, executives tend to exhibit more risk-averse tendencies.

ACHIEVING ALIGNMENT

With “an acceptable level of loss exposure” defined, we now have something to achieve alignment with. This subgoal is relatively straightforward, and has three subgoals of its own:

- Good risk landscape visibility
- Know the gap, and
- Close the gap

We will cover all three of these below, as well as some of the metrics that support them.

Good risk landscape visibility

We’ve already discussed visibility a bit, but here we’d like to make a not-so-subtle claim: it is impossible to achieve an acceptable level of loss exposure if your visibility into the risk landscape has gaping holes. Well, maybe not impossible, but it would be pure coincidence if an organization actually achieved alignment with its risk appetite while half blind. As a result, the first subgoal under achievement has to be to attain good visibility into its risk landscape, which requires metrics regarding assets, threats, controls, and impact factors.

Asset visibility

Common asset metrics tend to include how many assets there are, where they are located, what organization processes they support, who is responsible for them, etc. This includes system assets, network assets, application assets, data assets, and facilities, etc. Without this information, you can’t accomplish some of the other key elements in risk management, like risk assessments and analysis, with nearly as much confidence. No organization will ever have perfect visibility, but if visibility isn’t an explicit objective, there is a much higher probability of gaping holes. Organizations that have a decent handle on asset metrics tend to have mature asset management processes surrounding the introduction of new assets, changes to assets, and disposal of assets. Where this exists, it’s often part of the business continuity program, which may include a configuration management database (CMDB) solution.

Threat visibility

Threat metrics should, unsurprisingly from a FAIR perspective, focus on threat event frequency (TEF) and threat capability. For some threat communities (e.g., insiders of one sort or another), you can also include a metric regarding the number of threat agents, because there is likely to be some correlation between the number of threat agents and the probability of threat events (malicious or not).

Very few organizations really seem to leverage threat metrics. Oh, you'll often see things about the number of viruses blocked, the number of scans against web systems, and such, but beyond that, organizations tend to underutilize what could be a rich source of intelligence. Later in the book we give SIEM providers a hard time for not leveraging their data very effectively. Today nobody is asking them to be very proficient because common practices regarding threat metrics are usually pretty superficial. If you adopt FAIR as a fundamental component of your organization's risk management practices, you will inherently evolve your approach to threat metrics.

Control visibility

Control metrics are going to be primarily about the frequency, severity, and duration of variance. As you will recall from the controls chapter, variance is the enemy because variance from an intended state of control, almost always exists when a significant event occurs. In the controls chapter, we discussed that variance duration is comprised of two elements: the time it takes to discover a variance, and the time it takes to remediate variance. Also recall from the controls chapter, this exposure window data can be combined with TEF data to help us understand how quickly we need to remediate weaknesses.

Impact factors

There are three categories of impact factor metrics: asset characteristics such as volume; sensitivity and criticality; organizational conditions like stock price, cost of capital, compliance requirements, etc.; and external conditions like regulatory penalties. For the most part, these metrics are used within risk analyses themselves, rather than as a separately reported metric.

Visibility analysis

An approach we've used to gauge and manage visibility involves carving up the risk landscape into logical categories and then estimating the level of visibility the organization has into each category. For example, we might carve the landscape into: servers, web applications, network devices, databases, etc. For each of these, we would estimate the current asset, threat, control, and impact visibility on a scale of 0–100%. In other words, we would estimate how much we think we know about each metric dimension (asset, threat, control, and impact) of each category. If we believe we have perfect visibility into the asset dimension of servers, for example, our estimate would be 100%. If we weren't sure whether we even had servers in our environment, or we suspected we did but had no idea where they were, our estimate would be 0%. Yes, we know, we're asking you to estimate what you don't know, which some people really struggle with conceptually. There are a few things that can help with this, though:

- Just as we did when performing risk analyses, use calibrated ranges—e.g., "between 80 and 90%."
- Think about the processes, technologies, and other aspects of your environment that would affect this estimate. For example, if your organization has a lot of unmanaged databases and other "shadow IT" in it, then your visibility into those parts of your risk landscape is likely to be much lower.
- Another way of thinking about it is if tomorrow you suddenly gained perfect visibility into a category (e.g., servers), how far off would today's visibility estimate be? Were you off by 10% (i.e., your current visibility is 90% of what was actually out there)?

The point here, though, is not to have a perfect estimate. The point is to force the organization to think about visibility explicitly and differently than it has before, and then put into place the processes, technologies, etc. that allows it to improve its visibility. For example, let's say we did this visibility estimation drill for the databases, servers, and web applications in our organization and charted the results in Excel (see Figure 13.4). What can we tell by simply looking at the charts?

The vertical axis represents the estimated percentage of visibility (higher is better). Vertical box length represents the level of uncertainty in the estimate (smaller boxes are better).

First, it's pretty clear that there is good visibility in every dimension of the web apps. This may be because in this organization it's a smaller and more closely managed part of the landscape. The organization also seems to have a reasonably decent handle on assets in every category, perhaps because it has good asset management practices and change management. Likewise, it appears to have good controls visibility, perhaps because it has strong controls testing practices and technologies that provide up-to-date information on control conditions. For databases and servers though, it's clear that threat visibility isn't great, and there appears to be a lot of uncertainty about the estimates themselves. Also, there seems to be questionable visibility regarding database impact, which suggests that the organization may not know as much as it should about the sensitivity, criticality, or both, of these assets.

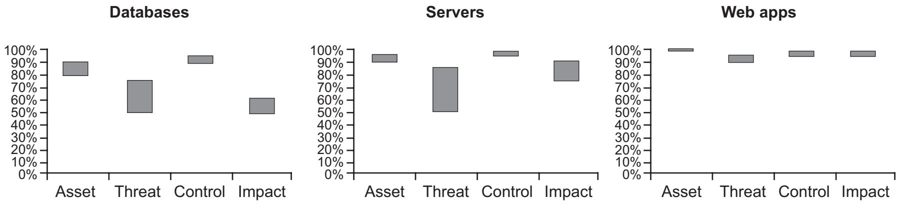

FIGURE 13.4

Visibility charts.

Having exercised this drill and come up with results like this, an organization can implement or improve processes and technologies that will improve and then maintain visibility. For example, maybe it needs to improve database and server logging and monitoring to improve threat visibility. Perhaps they need to review the purpose and content of databases to improve their understanding of sensitivity and criticality, and/or apply a classification scheme to them. Regardless, at least now they have a way to identify and address weaknesses in their visibility.

From a risk management perspective, an organization that has been through this exercise and acted to improve and maintain its visibility is obviously going to be much better positioned to achieve alignment with its risk appetite than an organization that hasn't.

Know the gap

You can't close a gap you aren't aware of. The good news, from a metrics perspective, is that we have already covered this one in the earlier discussion about establishing the organization's "acceptable level of loss exposure." Basically, the gap between the organization's current loss exposure and its acceptable level of loss exposure tells us how much further it has to go, which can drive the level of effort and resources an organization applies—i.e., the larger the gap, the more emphasis is going to be placed on closing it.

Close the gap

It probably goes without saying that achieving alignment with an organization's acceptable level of loss exposure most often involves risk reduction efforts. Certainly in some cases the coin will be turned and the organization will make decisions to increase exposure to loss, but this is the less common case in our experience. Regardless, there are two dimensions to this subgoal, and thus two questions the metrics need to help answer: "Where does risk exist?" and "What control deficiencies exist?"

From a level of effort perspective, the good news is that the same sources of data generally answer both questions. Risk assessments, self-identified reporting, and loss events are the three most common metric sources for this subgoal. Risk assessments and loss events are fairly obvious in terms of what they mean and the data they provide. Self-identified reporting, however, can take a couple of different forms, including personnel who notice that something is out of whack and report it, as well as tools or processes that are implemented specifically to identify and report on the condition of assets and controls (e.g., configuration monitoring tools, regular access privilege reviews, etc.). The process of closing the gap simply involves making changes in the people, policy, process, and/or technology that reduce the gap between the current state of loss exposure and the organization's risk appetite.

MAINTAINING ALIGNMENT

It doesn't matter whether we're talking about your bank account, your health, or the risk posture of your organization, maintaining the condition of anything over time boils down to preventing undesirable conditions, and then detecting and remediating undesirable conditions when they do occur. In order to accomplish this in risk management, you need to have visibility not only into the assets, threats, controls, and impact factors we discussed earlier, but also into the risk management elements (decision making and execution). When variance does occur, as it invariably will to some degree, you have to be able to figure out why it occurred so that you can address the root cause—particularly if it's a systemic problem.

We've already covered the first four metrics associated with this goal: asset, threat, control, and impact factor metrics, so we won't belabor them here. However the last two—decision visibility and execution visibility—warrant some discussion.

Decision visibility

Decision visibility metrics boil down to two subtypes; who is making decisions, and the information upon which they are basing their decisions. Data captured about decision-makers on their policy approvals, policy exceptions, change management approvals, etc. is important for two reasons: first, to be able to follow up on decisions that are questionable (either in who made them or why they were made), and second, to hold people accountable. The importance of this second point cannot be overstated. The level of scrutiny and caution many executives will put into a risk decision is different if there is a decent chance they'll have to explain their decision to someone up the food chain at a later date. Knowing which decision-makers are accepting the most risk, and thus which parts of the organization are introducing the most risk, can provide a useful picture of how risk is being managed.

Metrics regarding the quality of risk information are also important, and this is something about which very few organizations pay attention. Analyst training, qualification data, and metrics regarding analysis quality can be extremely useful. The table (Table 13.2) below provides a partial list of some of these metrics.

Table 13.2 Decision Visibility Metrics

Subcategory	Example Metrics
Analyses	Percentage of analysts certified in the organization's chosen risk analysis method
	Number of risk analyses performed
Analysis quality	Number of risk analysis reviews
	Number of risk analyses determined to be inaccurate
Decisions	Number of policy exception reviews
	Number of approved risk acceptances
	Number of high risk acceptances
Unauthorized decisions	Number of risk acceptance decisions by unauthorized personnel

TALKING ABOUT RISK

In one organization we've worked in, we managed to put a process in place that ensured information security was engaged to provide consulting and oversight for all IT projects. Just hard-wiring engagement like this into the IT project management process was considered a victory. At last, we would be able to really help ensure that the organization did not implement "risky stuff."

Based on our beliefs about previous project management practices and decision-making, we expected some percentage of projects to be slapped down by management for potentially introducing too much risk. What we encountered, however, was business as usual. Our infosec personnel would document their concerns during the course of projects and present their risk reports to management. Management would yawn and ask, "Where do I sign (to accept the risk)?" After a while, it became obvious that something wasn't right because a significant number of projects were being implemented where the level of security risk had been rated "high," and not one project had been rejected by management.

After digging into the metrics, we discovered that fully a third of the projects had been rated "high risk." This raised a concern because, as bad as we thought some common practices were, it seemed unlikely that the organization could have survived as long as it had if it was routinely implementing that much high risk into its technology landscape. A closer look at the specific findings behind those "high risk" ratings provided the answer—the majority of concerns that the information security personnel had been calling "high risk" did not pass the laugh test. For example, virtually any time a policy was being violated, it was automatically characterized as high risk. Surely, there are times when policy noncompliance represents high risk, but just as surely, this is not always the case. It's not even the case in most instances.

To resolve this problem, we first trained all of the personnel in a rudimentary form of FAIR analysis and gave them a simple process and set of matrices to use when determining the level of risk associated with any findings they might have. A few months later, we looked at the data again. The number of "high-risk" projects had dropped by more than two-thirds. In other words, fewer than 10% of projects now were labeled "high risk" (and there was still room for improvement, analysis-wise). At this point, we went to management, explained the new risk analysis approach, and advised them that if they saw a project labeled "high risk", they owed it to themselves to think long and hard about accepting the risk. A couple of months and a few more conversations later, executive management decreed that henceforth no projects would go into production with high risk unless acceptance was signed-off at the appropriate level; even then, there had to be a defined remediation plan with committed resources. In the following year, only one high-risk project went into production, and it had the appropriate approvals and remediation requirements.

Prior to doing decent risk analysis, management couldn't discern the important stuff from the "noise." Unfortunately, buried within that noise were conditions they did not want their signature next to—we just hadn't helped them recognize which conditions those were. After we improved our risk analysis methods, and could demonstrate it, their attitude changed and we finally started making a difference for the organization. Being able to use metrics regarding analyses and individual findings enabled us to recognize and address a critical problem that was affecting decision-making.

The bottom line is that if your organization doesn't at least occasionally examine the quality of the risk information being given to decision-makers, then you might be missing out on an important opportunity to improve decision-making. In a large organization especially, tracking the quality of analyses over time and across a population of personnel can help you identify analysts who need additional mentoring, or systemic problems in risk analysis and reporting.

Be particularly mindful of risk assessments performed by people outside of your organization. We regularly encounter third-party assessments that are embarrassingly bad in terms of how they rate risk. Left unchallenged, this informat ion at best creates unwarranted concerns. At worst, it could result in resources being redirected from more important objectives. We have regularly "sniff tested" third party findings (usually just the "highs" and "mediums") to see what percentage of them made sense and could be defended. It isn't unusual for us to find fewer than 50% that pass the sniff test.

TALKING ABOUT RISK

A few years ago I (Jack Jones) engaged a global security consulting practice to perform an attack-and-penetration exercise on the company I worked for as the CISO. Shortly into the engagement, the consultants approached me with some dire news. They had discovered several high-risk "vulnerabilities" in one of the most important corporate web applications, and were recommending aggressive remediation measures. I examined the findings and pushed back on the consultants. Yes, they had identified weaknesses, but had they considered the frequency of the kinds of attacks that would leverage those weaknesses? How about the frequency of any sort of attack against that application, and especially the part of the application where the weaknesses existed? How much skill was required to exploit those weaknesses? What kind of access to underlying sensitive data would be gained and/or what level of control over the underlying systems? After talking through these considerations, the consultants backpedaled and changed the "high severity" of their findings to "medium," and in several instances, to "low." As a result, my organization was able to appropriately prioritize its remediation efforts and avoid unnecessarily impacting key projects and business operations.

These consultants weren't stupid. On the contrary, they were simply following a script and using a method for rating risk that only looked at the control deficiency, and didn't think through the threat or impact aspects of the situation.

Execution visibility

These metrics help us to understand how the organization is doing in terms of managing the factors that affect execution. This usually includes things like policy and standards improvements (from a clarity, conciseness, and usability perspective), communications from management regarding expectations, awareness training, and enforcement (and rewards, if the organization is being smart about changing behaviors).

The table below (Table 13.3) provides some examples of execution visibility metrics that you might find useful.

Table 13.3 Execution Visibility Metrics

Example Metrics
Percentage of personnel who have had generic security awareness training
Percentage of personnel in specific roles who have received awareness training pertaining to their specific responsibilities
Number of personnel who were rewarded or commended for their security decisions and actions
Number of personnel who were reprimanded or fired for their security decisions and actions
The frequency of control reviews (broken out by technology, etc.)

Continued

Table 13.3 Execution Visibility Metrics—cont'd

Example Metrics
The percentage of assets where control reviews are taking place
The maximum, minimum, and mean time to discovery of a variance
The number or percentage of variances discovered through different sources (e.g., audit, security testing, self-reporting, third parties, incidents, etc.)
The maximum, minimum, and mean time to correction (may be broken out by control type, technology, department, etc.)

TALKING ABOUT RISK

In the past, we've had good success with leveraging metrics to generate healthy competition between business units related to audit findings, patching, and user awareness. Executives are often a competitive bunch, and this can drive focus and improved execution against key risk management objectives. It also can have a positive effect on the risk management culture of the organization. Like anything else, though, it shouldn't be overdone. Keep these competitions limited to issues that everyone agrees are important.

We also have set up processes to present individual awards to people (outside of information security) who demonstrated exceptional contributions to information security through their ideas for improvements or sharp attention to problems. People tended to take real pride in these awards.

REMEDIATE CAUSES OF VARIANCE

Where unintended variance exists, unintended risk exists. As a result, we need to be able to identify the existence of variance so that we can remediate it and, more importantly, with respect to maintaining a desired level of loss exposure, resolve the underlying causes of variance.

Variance data

Variance data come from the same sources we encountered for the close-the-gap goal: risk assessments, self-identified reporting, and loss events. From one perspective, the data are the same—they are the findings and deficiencies provided from those sources. From a metrics and decision-making perspective, however, the data take a different form. Instead of the "findings to be fixed" point of view in close-the-gap, variance metrics are intended to help us understand the frequency, severity, and duration of variance across the population of assets and controls. The difference boils down to the fact that the close-the-gap goal is all about, well, closing the gap, while the variance data is intended to help us manage risk over time by understanding and treating the sources of variance.

Examples of variance data are shown in Figure 13.5. In these charts, it appears that network systems are well managed, with little variance except for some patching. Databases and servers are struggling, however, particularly with patching (go figure).

You can (and usually should) get into much more detail with this kind of data, fleshing out things like time to discovery, time to remediate, severity levels, etc. For

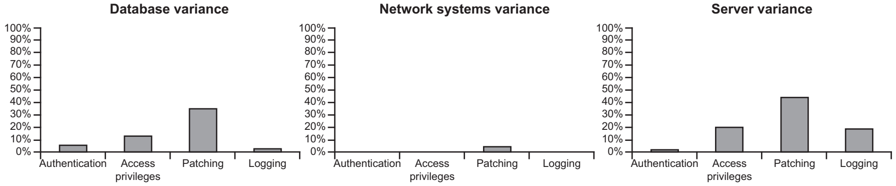

FIGURE 13.5

Variance charts.

example, even though databases may have a relatively high percentage of systems that aren't up-to-date with patches, the risk implications of those missing patches may not be as significant as another part of the environment that has a lower frequency of variance. The bottom line is, there is a lot of useful intelligence to be gained from these metrics about what is working and what is not working within your environment.

Root cause analysis

As we pointed out earlier, our profession is adept at identifying control deficiencies but not well practiced at figuring out why deficiencies come to exist, often repeatedly within an organization. In response to this problem, we developed a root cause analysis (RCA) process that some of you might think of as a prestructured "five why's" process for determining root cause. Not only does the process help identify the root cause for individual issues, it also can be used against a portfolio of issues to identify systemic sources of problems within the organization. Used in this way, very often an organization will discover that one or two fundamental problems are responsible for most of its control deficiencies. Resolving these systemic problems can have a significant positive effect on the organization's risk posture. We feel this topic warrants specific attention, so you'll find a detailed discussion regarding RCA in the next chapter.

TALKING ABOUT RISK

Deming had it right. If you aren't familiar with Edwards Deming, he is the engineering expert who helped the Japanese manufacturing industry evolve its practices and become the world's center of manufacturing quality excellence. A large part of Deming's focus was on managing variance in the manufacturing process and output. The result was better product reliability.

Well, information security is not that much different. As we've mentioned elsewhere in this book, bad things happen largely where variance exists. As a result, a significant part of our task in information security is to reduce the frequency, severity, and duration of variance. You can try to accomplish that by playing whack-a-mole and addressing each instance of variance individually, or you can seek out and resolve the fundamental sources of variance. We strongly prefer the latter.

Examples of the metrics that result from a root cause analysis are shown in Figure 13.6. As you can see from the first chart, the primary cause of variance was due to choices made by personnel. An analysis of those choices in the second

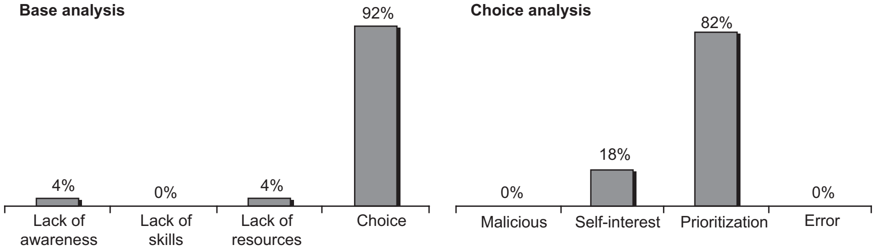

FIGURE 13.6

Rca metrics.

chart found that the vast majority was due to people prioritizing other business imperatives over information security. Additional analysis might find that these prioritizations were inappropriate and due to poor communication or enforcement by management, or because stronger incentives existed for the other choices (e.g., meeting deadlines and budgets).

COST-EFFECTIVENESS

We have said it before, but it bears repeating. Our responsibility as information security and risk management professionals is not to help our organizations manage risk. Our responsibility is to help our organizations manage risk cost-effectively. In many ways, the information security profession implicitly recognizes this and in large part tends to operate in ways that support this goal. In order to do this well, however, it needs to be an explicit goal with metrics that we actually pay attention to.

Not over or under control

Intuitively (or in some cases because business stakeholders have beaten this into our heads) we understand that there is a balance between risk, opportunity, and cost. Finding that balance inherently requires the ability to measure risk in a manner that allows us to compare, in like terms, where we are against where we want to be. We would not recite again the challenges our profession has faced to-date in this regard, but at least now we have models, tools, and methods that improve the odds of striking this balance. From a metrics perspective, we've already covered the two elements in this subgoal: an articulation of what constitutes an acceptable level of risk for the organization, and a means of measuring the organization's current level of risk.

Optimize solution selection

This is a sore point for us because we see relatively little attention paid to it in our profession, and it can be an opportunity to make a real difference for organizations. The archenemy of this opportunity is commonly known as "best practices." Don't get us wrong. There is a role for "best practices" in the industry as a point of reference, or if your organization is beginning from scratch and simply needs a place to

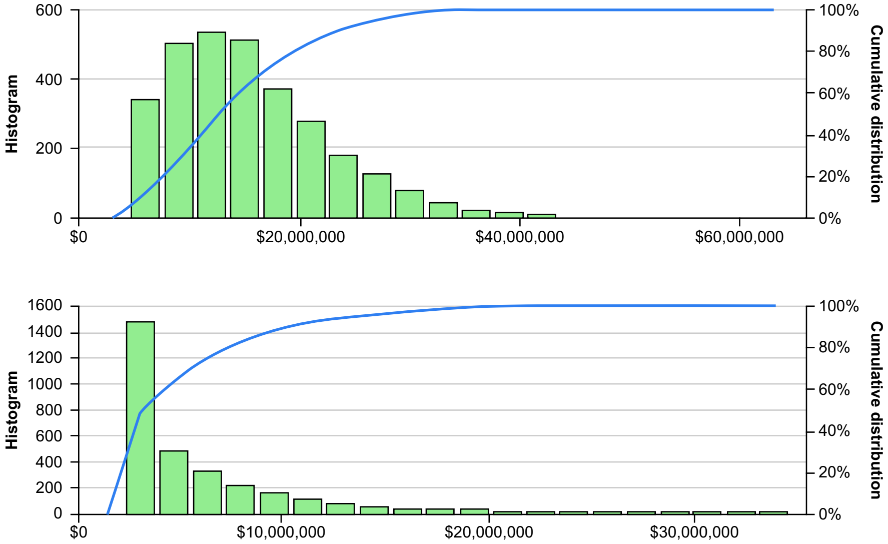

FIGURE 13.7

Loss exposure before and after controls.

start. What they should not be is a crutch, or a stick for beating management into doing what the information security organization, audit department, or regulators think it should do. Unfortunately, that is a common-use case.

TALKING ABOUT RISK

Sometimes you will hear people differentiate between "best practices" and "common practices." The differentiation is relatively common sense in that common practices (what the majority are doing) aren't necessarily the same thing as more advanced "best" practices. Very often however, the usage of these terms tends to get blurred. The bottom line is, if your organization wants to align with industry-adopted practices, it would be a good idea to decide whether it wants to align with common practices or best practices, or a mix in different parts of its risk landscape. Regardless of where and how it chooses to align, we strongly believe that those practices be considered guidelines, and that the organization should still exercise critical thinking to find "best fit" solutions whenever possible.

Being able to evaluate current state loss-exposure levels and the effect new or changed controls are likely to have on those levels can be incredibly powerful for guiding decisions. The charts in Figure 13.7 show the results of an analysis for current state loss exposure (top chart) and forecast loss exposure (bottom chart) after implementing a proposed control. This represents the benefit side of the cost-benefit equation. In order to do this, however, you have to be able to identify the control opportunities and estimate their effects. This is where FAIR and the controls ontologies can be leveraged to great effect. Also, the better your risk-related metrics, the more you can do in this regard.

The cost side of the equation is most commonly focused on the capital and operational costs of controls. Depending on the control, this may be all you need to focus on. In some cases, however, you also have to consider the opportunity costs to the organization—i.e., resources applied to information security can't be applied to other organizational objectives, and/or the drag on business processes that some controls introduce.

Acceptable rate of progress

This is another aspect of risk management that we see some attention paid to, but often not by information security—it's executive management throttling the rate of progress through its purse strings. Generally, information security is working to reduce risk as fast as its resources allow, and often complaining that it can't move fast enough due to a lack of resources. Well what is "fast enough?"

TALKING ABOUT RISK

Here again, it is important to keep in mind that executive management is faced with growing the business, managing operational expenses, and managing risk of various types, information security being just one of those types. It's also important to keep in mind that it is executive management's risk appetite and risk management pace that matters and not our own. If we're unhappy with the pace, then maybe we haven't done a good enough job of communicating risk to them so that their level of concern increases and the purse strings open up. Maybe it's simply a matter of them having even bigger fish to fry, of which we are not aware.

Having metrics related to the gap between the current level of risk and the organization's risk appetite, the progress that has been made since previous checkpoints, and the progress that's forecast for the next checkpoint can be huge in terms of gaining and keeping management support. It also provides a way to forecast the effects of risk management initiatives so that management can better understand the value proposition, and thus make better-informed investment decisions.

A visual representation like Figure 13.8 provides an at-a-glance view of past progress, current state, and forecasted progress, given current or planned resources. From this, management can better decide whether they want to pick up the pace, stay on pace, or even slow down. You can also break this down further into different dimensions of the risk landscape that are of particular interest or concern to management (e.g., online systems, SOX, etc.) so that the pace can be controlled more granularly. We've found that illustrating past progress is very helpful so that executive management can understand what they have gotten from their previous investments. Be aware though, that a postanalysis of previous risk management investments might show that they were not well chosen or executed, which could be disappointing news to your stakeholders.

Efficient operations

Operational efficiency has two primary questions associated with it. Both are important, but we would argue that if you only focus on one of these, focus on the first question.

- Is the organization focused on the most important issues and opportunities?
- Is the organization fully realizing the cost-benefit from its resources?

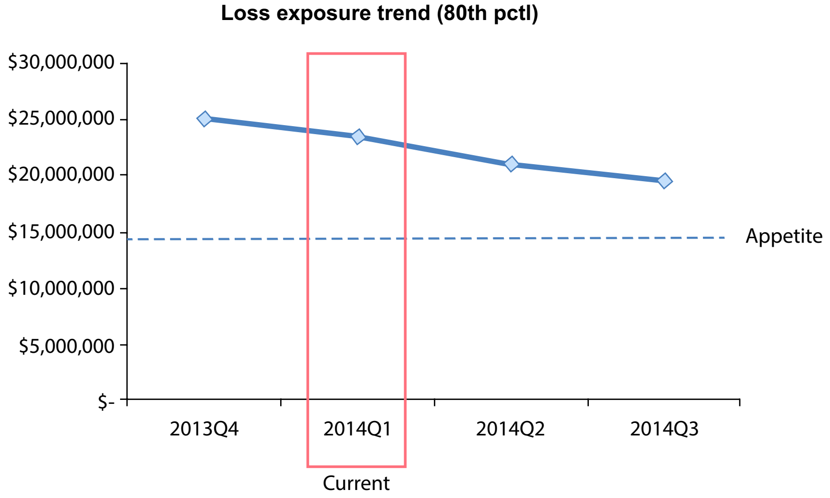

FIGURE 13.8

Loss exposure to tolerance over time.

Focusing on focus

You may have noticed that so far we haven't talked about prioritization. Identifying and fixing problems, yes, but not in what order. Your wait is over! This is, of course, one of the most important elements in risk management and a prime opportunity for metrics. Unfortunately, it's also where our profession has struggled mightily.

The focus question has two categories of metrics to it:

- Metrics that help us to understand where risk is concentrated, and
- Metrics that help us to prioritize the conditions that provide the best risk reduction and/or risk management bang-for-the buck

The first metric comes from examining the risk landscape and being able to measure the level of risk in each area of that landscape. Being able to do this, however, requires that we carve up the landscape in some logical manner. Maybe it makes sense in your organization to parse the landscape from a technology point of view—e.g., databases, servers, personal systems, applications, network systems, etc. We have seen this work well. We've also seen organizations carve the landscape up by business process—e.g., online retail, finance, HR, research and development, etc. Another option that we kind of like is a mixture. For example, an organization might identify a handful of key business processes as distinct analytic targets, and then carve the rest of the landscape up by technology category. From an analytic perspective, it doesn't matter that much. It's whatever fits best for your organization.

Performing a risk analysis on each of these landscape elements allows you to compare them and recognize where the organization has the most risk and thus should focus its efforts. Figure 13.9 shows partial results from one such analysis.

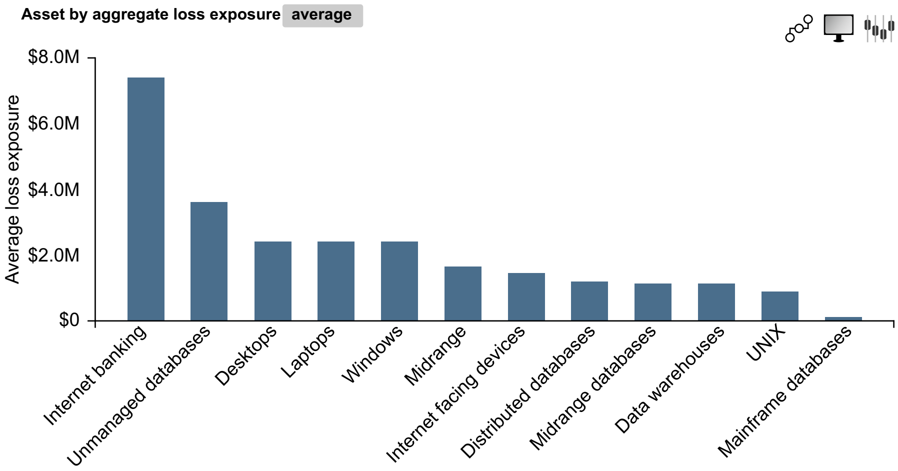

FIGURE 13.9

Asset by aggregate loss-exposure.

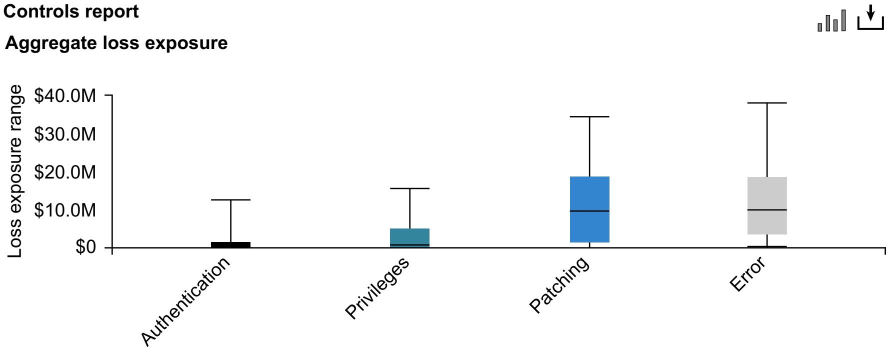

FIGURE 13.10

Aggregate loss-exposure by control.

The second metric (Figure 13.10) comes from examining the elements that are driving risk into the organization; typically, deficient controls. By understanding which controls are responsible for the most risk, particularly within the highest areas of risk concentration, you're able to manage risk much more cost-effectively.

By extending this analysis into the variance and decision-management controls, you're able to not only remediate risk cost-effectively, but also improve risk management functions cost-effectively. This type of analysis and reporting takes us much farther along the maturity continuum than where we've been with the standard risk assessment containing a list of "risks" that are really control deficiencies.

Leveraging resources

This last metric component is another underutilized opportunity in most organizations. The first part has to do with fully utilizing our personnel. Do we understand what their skills are, and are we effectively leveraging those skills? From a metrics perspective, we can survey our information security personnel for their backgrounds, skills, education, certifications, and even personal preferences. With this information, we can better place people where they will provide the most value to the organization and have the greatest likelihood of personal satisfaction and professional success.

TALKING ABOUT RISK

As a new leader in an organization, one of the first things we do is meet with each person in our organization to learn about their backgrounds, skills, work preferences, etc. Every time, this has resulted in significant changes to work assignments, and in some cases, the organization's structure. Every time, morale and productivity also improved significantly.

The second part of this metric component is maximizing the utility from tools the organization already has in place or chooses to deploy. It involves mapping the capabilities of tools against the controls ontologies so that organizations can more fully recognize and leverage the value of its investments. Once the capabilities are mapped, an organization can overlay that mapping against its risk landscape to identify where strengths and weaknesses exist.

TALKING ABOUT RISK

FAIR and the controls ontologies also can be a useful way of evaluating the value proposition of vendor products and services. In fact, in the past we've been asked by vendors to evaluate their products' value propositions so they can improve their marketing message. We've been happy to do this, but sometimes the results aren't what the vendor expects…

THE FEEDBACK PERSPECTIVE

Do you recall the diagram in Figure 13.11 from the chapter on risk management, and how in order to manage the risk landscape as a system, you need to use metrics as feedback?

Well, if we were to map each metrics defined from our GQM outline to this diagram, we would find that every element in the landscape is covered. What this means is that by following our GQM framework, an organization is naturally managing the risk landscape as a system.

MISSED OPPORTUNITIES

One of the things that baffles us is when people in our industry say, "We don't have enough data." In almost every case, this couldn't be further from the truth. In fact, with all the tools at our disposal these days, we are practically awash in data. As a profession,

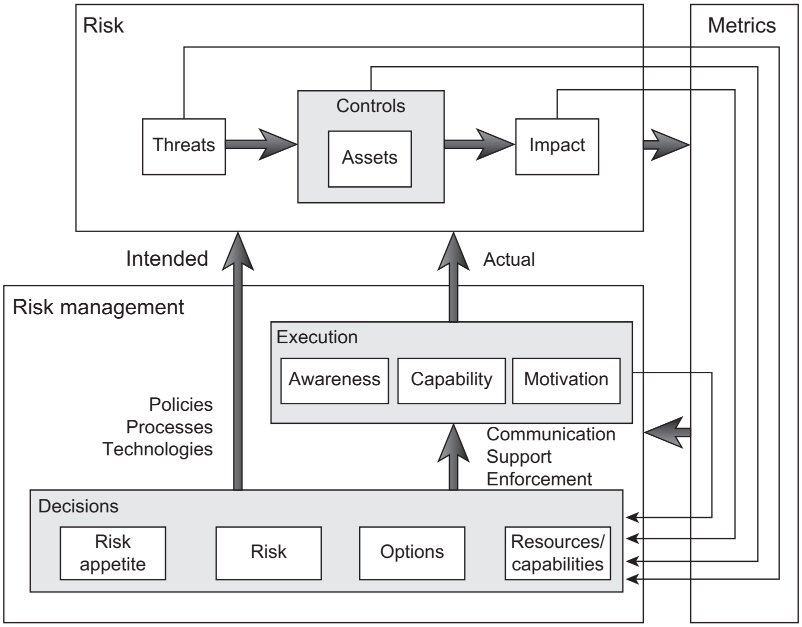

FIGURE 13.11

Risk management.

however, we haven't leveraged the available data effectively, in large part (we believe) because the profession hasn't had the models to tell us which data we needed to answer which questions. With the FAIR ontologies and the GQM framework discussed earlier, this shouldn't be nearly as significant a problem going forward. In other cases, we simply fail to recognize some of the metrics opportunities staring us in the face, and yet at other times, the data our tools provide us fall well short of their potential.

LOSS METRICS

The silver lining to loss events is that they can be a marvelous source of information to improve the quality of risk analyses and our understanding of the risk landscape. Unfortunately, we encounter very few organizations that fully leverage this data. As you review the information below, ask yourself these questions:

- How fully does your organization take advantage of its incident data?
- What would it need to do to improve in this manner?
- If your organization fully leveraged its incident data, how much would it help you when doing risk analysis?

The answers to these questions for most organizations are "*not very*," "*some relatively simple process and documentation changes*," and "*a lot*." The point is that many loss metrics aren't difficult to acquire, and they can offer significant value in helping us understand our risk landscape.

There are several categories of loss metrics that we like to use:

- Incident counts, which tell us about loss frequency
- Loss magnitude data, which helps us understand incident impact
- Threat data, which tells us who or what caused the incidents (including motivation)
- Control data, which tells us what failed (and hopefully also what worked)

You'll notice that these metrics are more than just how often bad things happened and how bad they were. They also have to do with answering the question, "why?" The table below (Table 13.4) provides some examples of these metrics.

Most organizations capture a fraction of this information, yet it's all there to be had. Does it take additional work? Sure it does, but once capturing these variables becomes an inherent part of the incident management process, the relatively little additional work is well worthwhile.

FALLING SHORT

First of all, in the interest of full disclosure, we either currently or in the past have worked for vendors, so we've been as guilty as anyone else in the deficiencies we are about to describe. Hopefully, by discussing the weaknesses we see in vendor metrics, the bar will rise a bit and our profession will begin to see meaningful improvements.

If you work for an information security vendor, ask yourself whether your products or services can improve in their ability to provide meaningful metrics, particularly given what has been covered in this book. With this in mind, we're going to take this opportunity to use the SIEM vendors as an example and suggest that they are falling short of their potential from a metrics perspective. From what we've seen and experienced (and admittedly we haven't seen them all), these products and services largely focus on letting their clients know that a particular event is happening or nothing wrong with this. has happened. In other words, they're primarily an incident detection and reporting service. Sometimes they also include data regarding the types of vulnerabilities on systems they monitor. There's nothing wrong with this. It's very important data. What they aren't doing, however, is fully leveraging the treasure trove of data at their disposal.

For example, as a CISO we would love to know not just that web application X has been subjected to 20 attacks in the past month, but that 18 of those attacks also were directed against every other application on our network perimeter. Furthermore, of those 18 attacks, 15 of them were also thrown at other unrelated businesses on the Internet at roughly the same time. What does this tell us? It suggests that of the 20 attacks, five were targeted at our organization, and two of those were targeted specifically against application X.

Table 13.4 Loss Metrics

Category	Example Metrics	Discussion
Incident counts	Malware infections Phishing compromises Lost/stolen laptops, etc. Applications compromised	These metrics are important but do not require much explanation. The more interesting loss metrics have to do loss magnitude, threats, and controls.
Loss magnitude data	Number of records lost	How many sensitive records, if any, were compromised? This can, and should be, further broken out by type (e.g., social security numbers, credit card numbers, other PII, corporate sensitive information, etc.).
	Number of customers notified	This is relatively self-explanatory. If we also know the per-customer notification cost, then this information helps an organization pin down response costs from events.
	Credit monitoring offered	How many of the people whose information was compromised were offered credit monitoring?
	Credit monitoring accepted	How many of the people who were offered credit monitoring actually took advantage of the offer?
	Customer churn	If an event involved customers (or business partners for that matter), how many of them took their business elsewhere? This data is rarely documented or explored after an event, but it is there to be had and can really help an organization to understand the sensitivity of its customer base to events of various types. A related piece of information has to do with the lifetime value (in terms of profit) of customers who leave. Note that this does not have to only occur after a confidentiality breach, as severe outages or data integrity issues can also drive customers to leave.
	Fines and judgments	The good news is that most organizations do not have events that involve fines and judgments. But when they do, this is critical information to capture.
	Legal expenses incurred	Regardless of whether fines or judgments occur, it is not uncommon for organizations to engage internal or external legal expertise. When it is internal legal resources, these are “soft monetary” costs in terms of person-hours. For external legal resources, it is cold, hard cash.
	SLAs missed	Not only do we need to know whether a service level agreement was missed, but also the effect. Did the organization have to discount some invoice or refund some money?
	Public relations costs incurred	Following particularly nasty events, organizations may end up spending significant sums on managing their reputations. These costs should be captured when they occur.

Forensic expenses incurred	In many incidents, it is necessary to apply forensics to understand exactly what happened. When this occurs, whether using internal resources or external consultants, the costs should be captured.
Number of personnel affected	When there is an availability outage or significant degradation in system/application performance, it is important to capture the breadth and severity of the effect on personnel. Combined with an average, loaded hourly wage rate, this can help an organization gauge the productivity impact of events.
Person-hours expended in response	Response costs include the person-hours expended in dealing with an incident. Knowing how many people are involved, for how long, and at what loaded hourly wage, can help an organization to understand these costs.
System down-time	Self-explanatory.
Lost revenue	Sometimes when business processes are affected by outages, revenue loss occurs. A common practice is to use the length of an outage times the average revenue generated during that time of day and/or year to derive lost revenue. That often overlooks the fact that many times revenue is delayed rather than lost. In other words, depending on your business model, competition, etc., your customers may simply wait for the crisis to pass and then resume their business without an actual loss of revenue. Consequently, in order to get this metric right, you need to differentiate between lost versus delayed revenue.
Affect on stock price	Yes, we know this one is tough for a lot of reasons. Furthermore, few organizations to-date have had events been bad enough to demonstrate reputation damage as an effect on stock price. If your (publicly traded) organization suffers a significant event, however, you would be remiss to not pay very close attention to the effect on stock price. NOTE: A close cousin to stock price reduction is an increase in the cost of capital. There are now a few organizations with incidents so severe that they have started paying attention to this.
Insurance paid	With the increased use of cyber insurance, another metric worth paying attention to is how much (and what types) of the loss were covered by insurance.
Insurance costs	Assuming that your cyber insurance provider had to satisfy a claim your organization made, it would be good to know whether the insurance premium increased and, if so, by how much.
Fraud losses	Self-explanatory.
Fraud loss recovery	When fraud occurs, especially online fraud, there are often opportunities to recover the funds. This metric helps us to measure the efficacy of our recovery capabilities and thus our actual fraud losses.

Continued

Table 13.4 Loss Metrics—cont'd

Category	Example Metrics	Discussion
	Time to discovery	A critical piece of information is how long it took us to recognize that the loss event occurred. This not only informs us about the strengths and weaknesses of our detection capabilities, it also helps us to estimate the value proposition of improved detection capabilities—e.g., how much less loss would have occurred if we had detected the event sooner?
	Time to contain	After an event has occurred, how long did it take us to contain the event—e.g., bring up backup systems, kick the hacker off our systems, etc.?
	Time to resolve	How many hours, days, weeks, or months did it take us to resolve the event and return to "normal?" Normal often being defined as that state in which the event is no longer the topic of specific tracking and management attention. This can be anything from hours for simple outages to months (many months) for some catastrophic events.
	Remediation cost	Sometimes we have to implement additional technologies, policies, people, or processes in order to resolve an event. Documenting those costs is also important in order to understand the true effect of an event.
Threat data	Threat actors	Documenting who or what the actors were (or who were believed to be) helps us to understand what we are up against.
	Threat capability	Not all threat actors are created equal, nor do all compromises require rocket science to execute. By evaluating the level of expertise that was required to create a compromise, we implicitly gauge the efficacy of our controls. This can be incredibly informative in evaluating our loss exposure going forward and in making business cases for change.
	Threat intent	This is relatively straightforward. Was the actor looking for financial gain, to drive an ideological agenda, or just to hurt the organization (or was it Mother Nature simply acting out)? Of course, unless the actor is known, this may be somewhat speculative and based on inference given the type of event and the specific actions taken by the actor.
Control data	Control failure	This simply requires us to document what type of control (or controls) failed that resulted in the loss event materializing. Was a password compromised? Was a system not up-to-date on patching? did anti-malware fail to deal with a new virus? Did employment screening fail to detect a criminal history? Was someone untrained on how to recognize a phishing attempt? Did someone fail to label a network connection properly and as a result someone else disconnected the wrong cable?
	Control success	This is information helps us to understand which controls are helping us to limit the severity of loss events that occur. Unfortunately, this is commonly overlooked and undocumented. Did someone recognize that another employee was engaged in illicit activity? Did someone breach an account but was prevented from accessing sensitive information because access privileges were properly constrained, or because a system configuration setting limited access?

Going a step further, we would love to know more about the nature of those five attacks. How sophisticated did they appear to be—i.e., where did they fall along the threat capability continuum? Were they automated or manual? What types of weaknesses did they probe for? How persistent have they been—e.g., have they occurred in the past from the same source? Is that source focusing solely on our organization or do they also appear to be targeting our competitors? This kind of richer intelligence can be incredibly meaningful to an information security program that's paying attention. All of a sudden it's much easier to identify and respond to subtle changes that previously weren't discernible.

We also regularly encounter products that use significantly flawed models (yes, all models are inherently flawed to some degree, even FAIR) to drive both the data they collect and the analytics they apply to the data. In our experience, the products historically most challenged in this regard are vulnerability scanners. Very often, these tools use analytic functions that have relatively serious flaws, such as:

- Performing somewhat questionable math on ordinal scales
- Weighted values that don't appear to be substantiated in any clear fashion
- Variables applied to the wrong part of the equation, and
- Risk equations that are missing key variables

That said, some of the variables those tools capture can be extremely useful in our analytics, but we never rely on the risk scoring these tools provide because of our concerns. The good news is that two things are taking place as this book is being written: (1) a new version of CVSS is coming out, which we've been told has some important improvements; and (2) we're working with a vulnerability scanning vendor to improve their models and use of data. The results should be a much more highly refined and accurate output that will help their clients prioritize remediation much more effectively and efficiently.

The bottom line is, many security products seem to be missing the metrics boat to a greater or lesser degree. With any luck, some of what we've covered in this book will help.

KEY RISK INDICATORS AND KEY PERFORMANCE INDICATORS

Key Risk Indicators (KRIs) and Key Performance Indicators (KPIs) can play a critical role in managing risk effectively, or they can waste an organization's time and resources. Consequently, it's a good idea to get them right. Before we dive into the details though, we'd like to share some ground rules regarding KRIs and KPIs so that we're all on the same page:

- The word "key" is, well, key. The challenge, of course, is where to draw the line between those metrics that qualify as "key" versus those that don't. As far as we can tell, there is no clear, bright line to guide us, so each organization will have to make its own determination. Below, we'll discuss some of what, in our minds, are considerations that make a metric "key."

- "Risk indicators" should be metrics that inform us about how much loss exposure we have right now or how it's trending. The point is, KRIs are about *risk*. You know—threat conditions, asset-level control conditions, and the assets themselves.
- "Performance indicators" should be *execution and decision-making* metrics informing us about the risk management elements that contribute to our current level of risk and that affect our future level of risk.

Another consideration to keep in mind regarding KRIs and KPIs, is who the stakeholders are that will be using these metrics and what decisions are potentially going to be affected. For example, only a subset of the CISO's key metrics is going to be directly relevant to the Board of Directors and senior executives. In large part, the difference has to do with temporality and level of detail. The CISO will need to recognize specific risk or performance conditions that can have more immediate effects on loss exposure and that may require more immediate changes to processes, resource allocations, or technology (or the need to report something up the chain). Senior executives, on the other hand, tend to be more interested in whether things are or are not on track relative to earlier decisions and priorities, and whether they need to be particularly worried overall.

For that reason, we differentiate between what we refer to as strategic versus operational KRIs/KPIs. We'll spend most of our time discussing the operational KRIs/KPIs because they are often what feed the strategic ones. That being the case, you can assume the following sections are talking about operational metrics unless we specifically say otherwise.

Clues to "keyness"

KRIs and KPIs exist solely for the purpose of telling us that something important, from a risk or performance perspective, is either amiss or trending toward being amiss so that we can make adjustments and avoid unacceptable outcomes. This basic description tells us two things we need to do in order to recognize which metrics are key: (1) identify what's important to us, and (2) identify the variables that contribute to whether something has become, or is becoming, a problem.

Importance

Importance tends to differ between KRIs and KPIs. With KRIs, importance is based directly on asset value/liability—i.e., potential negative impact. Because KPIs are focused on decisions and execution that have some future affect on risk, importance is based on how quickly and profoundly that affect can take place. In either case, it boils down to assets at risk.

When considering assets at risk, identifying what is important should be relatively straightforward, at least conceptually. They are those things that represent the greatest potential for harm if negatively affected. In practice though, it can be a bit more challenging. One of the aspects that can make it more challenging is the notion of intrinsic asset value versus inherited asset value. For example, a bank safe has intrinsic value associated with how much it cost to purchase and install. The inherited value of the safe, however, includes the value of everything it protects. With that in mind, the more an asset contains, the more value it has. For example, the intrinsic value of a password is low, but the inherited value may be high depending on the types of information and system assets it protects. Likewise, an Internet-facing firewall may have nominal intrinsic value, but its inherited value tends to be extremely high because of everything it helps protect. (For the programmers in the audience, this resembles an object-oriented mode of thinking).

A last consideration regarding inherited value is that sometimes containment is not obvious. For example, if two servers have a trust relationship with one another from an authentication and access privilege perspective, each server inherits the value of the other because a breach of one represents a potential breach of the other. As you can see, this would become very complicated, very fast, if you tried to tackle it at a fine level of granularity. Fortunately, we can abstract the problem to a coarser level of granularity by essentially considering those two servers to be one “virtual” asset. This coarser granularity sometimes comes at the cost of precision in terms of estimates regarding control conditions, etc., but it should not affect accuracy, and it is very often a worthwhile tradeoff in terms of reduced complexity.

When trying to identify the importance of information, there are four primary considerations: (1) the value (or criticality, if you prefer) of the information in terms of helping the organization meets its objectives, (2) the potential liability associated with the information, (3) the volume of information, and (4) the cost associated with reconstructing the information if it is lost or corrupted. Note that confidentiality, availability, and integrity each can affect information differently.

When considering physical assets, importance boils down to two things: (1) the cost to replace the asset, and (2) the criticality of the asset in terms of its role in helping the organization meet its objectives.

A final thing to be aware of is that the “keyness” of some metrics may vary over time. Although most are likely to remain “key” indefinitely, others may come and go from the list of key metrics within an organization. For example, if an organization has historically struggled to comply with a specific policy expectation (e.g., access privilege management) *and* as a result represents a level of risk that has gained executive attention, then metrics regarding that condition may become key because it has been identified as something of particular interest to executives. Once the necessary improvements have been made and maintained long enough to demonstrate stability, these metrics may no longer qualify as key. Another example of a temporary key metric might be the achievement of project milestones on a large initiative that has executive leadership interest.

Condition factors

Having laid the groundwork for the “importance” aspect of KRI/KPI development, it’s now time to consider the factors that contribute to whether the condition of important assets is acceptable or not.

More about KRIs and KPIs

In this section, we will provide some high-level examples of the kinds of metrics we focus on for KRIs and KPIs. We'll also discuss why we focus on these, and share some considerations that can help guide you in developing your own. Please keep in mind that your needs may be different, and thus your KRIs and KPIs may be very different from the ones presented here.

Digging into KRIs

We carve KRI metrics into four categories: loss metrics, TEF metrics, threat capability metrics, and control condition metrics. We'll discuss each of these below. Before we do, we'll reiterate that KRIs are about loss that has materialized or that the organization is exposed to.

Loss Loss metrics, particularly loss magnitude numbers, are often considered to be of more interest to senior executives than CISOs. That said, a significant jump in the frequency of loss events, even when loss magnitude is low, can reflect changes requiring immediate attention. For example, if an organization detects a significantly higher number of malware that have made it past perimeter defenses (these are loss events even if the loss magnitudes are often nominal), it may point to an increase in threat capability, a decrease in control efficacy, or both. Similarly, an increase in the frequency or severity of loss events associated with new software releases might indicate problems with developer skills, testing processes, or project deadlines. The good news is that loss events are a great source of risk intelligence. The bad news is that they are lagging indicators that suggest our leading indicators need some work. Even that is good news in a way, though, because at least we're in a position to recognize and deal with it.

Threat event frequency In order to qualify as "key," you will generally want these metrics to focus on especially sensitive or critical assets or locations within your landscape. Examples include critical websites, large data stores, and assets with high levels of inherited value (e.g., the network perimeter). At the time of this writing, threat agents are increasingly focused on attacking users and end-user systems, which means TEF metrics at those points of attack are becoming more important.

Once your organization has established what the "normal" level of threat activity looks like against these points of attack, what you are looking for is to be able to recognize significant changes in threat activity. We're less interested in the fact that we had 50 attacks against our key web applications than we are that this represents a 50% increase over the previous week. Unusually high levels of threat activity are of obvious importance and may drive the need to alter controls. Unusually low threat activity can be equally important and may mean a couple of things:

- Threat actors have shifted their focus to another part of your environment (or to some other organization entirely),

- Your detective controls are no longer working due to unintended changes in those controls, or threat capabilities have evolved in a manner that allows them to evade detection.

Trends regarding shifts in TEF can be considered operational or strategic, depending on time horizon. This is an example of where good threat intelligence, typically from outside the organization, can be especially helpful.

Threat capability It's important to recognize when significant changes occur in the sophistication of the threat landscape. With regard to malicious actors, we are first and foremost interested in changes in the sophistication of attacks being experienced by the organization or by other organizations in the same industry. Secondarily, we also want to know when capabilities have changed regarding important technologies used within the organization (e.g., new zero-day exploits against Oracle databases if the organization uses Oracle for key business processes or to manage particularly sensitive information). We do not, however, care much about changes in capabilities that aren't relevant to the organization. For example, if the organization doesn't use a particular flavor of UNIX then metrics regarding new zero-day exploits of that variety are just noise.

As with TEF, intelligence regarding trends in threat capabilities may be operational or strategic, depending on the timeline. In our experience, about the only place to get good leading intelligence regarding changes in threat capability is from commercial threat intelligence providers because they have the resources to do this well (breadth of vantage point and depth of specialization by their staff). This can be expensive though, so you'll want to be sure to evaluate its value against your other opportunities.

From a nonmalicious perspective, sometimes a sudden, large turnover in staff or other events can represent a change in threat capability. For example, as you may remember from the analysis chapter, you can think of every new software release as a threat event (because new releases sometimes result in loss). A significant turnover in development staff could represent a lowered capability, and thus an increased likelihood of losses resulting from faulty software.

Control conditions The first criteria for identifying control KRIs is that only asset-level controls qualify. In other words, variance management and decision management controls will never qualify as KRIs. KPIs yes, but not KRIs. The second rule is that, as we discussed earlier, only those controls applied to particularly important assets should qualify as key. For example, if we've identified a particular set of databases as key, then we'll want to maintain a clear understanding of the resistive, detective, and responsive controls surrounding them.

For the most part, the characteristics we're interested in for these controls are variance frequency, severity, and duration. These metrics help us to understand the level of loss exposure on key assets. There may be some instances however where the intended state of control as defined by policy or standard contributes to an unacceptable level of risk (e.g., an encryption standard that hasn't changed in years and is no longer strong enough to resist common attack methods).

TALKING ABOUT RISK

When we focus on variance, there is an inherent assumption in play that the intended state of these controls (the expectations set through policies and standards) is sufficient to meet the organization's risk needs. For example, if an organization's policies or standards don't set an expectation that logs should collect information to enable event detection, or doesn't require monitoring of the logs, then the problem isn't variance from the organization's formally defined expectations. However, it might represent a variance from executive management's undefined expectations or from an external stakeholder perspective (e.g., regulators). Regardless, when policies and standards are deficient in some manner, that's a decision-management failure, which can be a function of poor visibility, poor risk analytics, poor communication, etc. These decision-making failures can be identified through the root cause analysis process we introduce in the next chapter.

Digging into KPIs

We like to carve KPIs into three categories: visibility, variance management, and decision-making. A good argument can be made that visibility falls into decision-making because it plays a big role in the quality of information decision-makers are operating from. That said, we believe it's an important enough concern to warrant its own KPI category.

Visibility Once an organization has identified its current level of risk landscape visibility, it is extremely important to quickly recognize when key parts of that landscape have changed. This boils down to changes in the things we want to have visibility into (e.g., new network connections, key websites, technologies, etc.) as well as changes in our ability to understand the threat landscape they face and the control conditions surrounding them. For example, CISOs may want a visibility KPI that tells them if important parts of their landscape have changed in some way (e.g., an extension to the network thru an acquisition, new large concentrations of sensitive information, etc.). Another KPI may be one that will let them know if visibility into key parts of the landscape has diminished (e.g., the threat intelligence service provider contract has expired, the DLP product we use to scan systems for sensitive information is on the fritz, or security testing of a key web application has been suspended).

Variance rates As discussed above, as a KRI these metrics fall within the control conditions category because they inform us on the level of risk that exists. As a KPI these metrics inform us about the effectiveness of decision and variance management (expectation setting, communication, support, and motivation). Because these metrics provide such effective optics in both directions—risk and risk management—it is perhaps our favorite metric. It's as close as you will come to the one metric that rules them all.

There will always be some amount of control variance across a set of assets over time. From a KPI perspective, being able to detect when variance increases in frequency, severity, or duration enables us to more rapidly identify and resolve the root causes of variance earlier rather than later, which helps keeps risk management on a more even keel.

Variance causes From an operational perspective, it is crucial that management is aware whenever variance is due to malicious intent or self-interest, regardless of the asset that's involved. Only then can they respond effectively in terms of managing the personnel or organization issues that set the stage for those choices. It also can be important to recognize broken policies, processes, or execution shortcomings that can affect key assets. Other than these, the causes of variance tend to be more a strategic metric.

An example of these metrics from a strategic perspective would include changes in the percentage of variances that are due to a lack of awareness, incorrect prioritization, lack of skills, etc. This information allows the organization to make the necessary systemic adjustments and stay out of groundhog day.

On-time closure rates This metric tends to be more strategic in nature than operational. What it boils down to is that the ability of an organization to consistently meet its commitments in terms of on-time closure of security findings (be they from audits or some other source) can be an important indicator of its overall commitment to risk management. Of course, it can also be an indication that (1) the organization is lousy at estimating how long it will take to remediate a problem (which could be a project management problem), and (2) the organization is lousy at measuring and/or communicating risk. Doing a poor job of measuring and communicating risk often manifests as low levels of commitment by management on remediation efforts. Regardless, if an organization consistently struggles to meet these commitments, some root cause analysis is in order.

High risk acceptances Most organizations will want to keep a very close eye on the rate of high-risk acceptances. This is a lot easier and more pragmatic to accomplish, of course, if the organization is performing good risk analysis (minimizing the signal to noise problem).

Unauthorized risk acceptances This metric can straddle the line between being operational or strategic. From an operational perspective, knowing where within the organization people are making unauthorized risk decisions can be critical if it's occurring around key assets. From a strategic perspective, this indicates that executive management may need to reinforce its position on compliance and risk-taking, or terminate some bad apples. Regardless, keeping tabs on how often personnel are making decisions they aren't authorized to make is pretty important. Also, if people know this is considered a KPI they are likely to be especially careful not to cross the line. To do this effectively does require very clear guidance on who is authorized to approve what.

Analysis quality Metrics regarding the quality of risk analysis are primarily strategic in nature, although they can have operational risk implications when risk decisions are being made regarding key assets. Note that this metric may be one of those that only qualifies as "key" until an organization has become comfortable that it is consistently being done appropriately. For most organizations today, we'd argue that is not the case.

The bottom line for KRIs and KPIs are that they are just the subset of the overall metrics we discussed earlier, focused on those parts of the landscape that we need to pay the closest attention to.

Thresholds

It may come to pass that some KRI or KPI metric has changed sufficiently that we need to do something about it or at least consider doing something about it. The notion of "changed sufficiently" (an alerting threshold, or trigger) is another parameter that has to be established. Most of the KRI/KPI thresholds we encounter are somewhat arbitrarily chosen. By that we mean when we ask people to explain why a certain threshold was chosen, the answers tend to be the equivalent of, “Because it seemed about right.” Actually, this is neither surprising nor necessarily bad because at the end of the day these thresholds are mostly about comfort levels—i.e., the point at which stakeholders are worried enough about an issue that they want to be prompted to decision or action. That said, if an organization has set an explicit risk appetite it can more effectively gauge the relevance of changes in KRI’s and KPI’s and thus calibrate its comfort levels.

DASHBOARDS

You may have noticed that we haven’t once uttered the “D” word in this chapter. That isn’t because we are “saving the best for last” or because we believe dashboards are inherently bad. No, we’ve waited until now because dashboards are simply tools for effectively and efficiently communicating the kinds of things we’ve been discussing throughout this book and especially in this Chapter. In fact, we would argue that you can’t effectively leverage dashboards unless you have a firm grasp on what we’ve covered above. Otherwise, what you can end up with is a bunch of charts and graphs that have limited utility and a low signal to noise ratio. Because there are other good references available on dashboards and data visualization this will be a small section. Besides, we think we’ve already armed you with much of the information you need to use dashboards well.

Dashboard content

Dashboards should be designed to meet specific information needs of specific audiences. That being the case, your first steps are to identify those audiences and their information needs. Not coincidentally, this aligns well with our discussion regarding KRIs and KPIs. That isn’t to say that dashboards should be limited to KRIs and KPIs, but that’s where we would start. This especially makes sense if dashboards are supposed to provide fast and simple access to your most important information. So if you know what your KRIs and KPIs are, and what decisions you expect to drive using the dashboard, then knowing what to include in your dashboards is a piece of cake.

TALKING ABOUT RISK

One of the challenges today is that virtually every information security tool seems to come with its own dashboard. As a result, you end up with a bunch of different dashboards, each containing specialized information regarding threat conditions, control conditions, assets, etc. What is missing is the overarching view—the one dashboard to rule them all—that takes all that disparate information and translates it into a single meaningful source of intelligence. Some of the GRC products are trying to play that role, and we’re beginning to see movement in this direction with other information security products, but from what we’ve encountered so far there is a lot of opportunity for improvement in this space.

With regard to a dashboard’s look and feel, as we have mentioned above, less is more. By that we mean that a dashboard should be clear and concise and present data in an uncluttered fashion. We are big fans of Dr. Edward Tufte and similar data visualization experts who advocate clean, simple, and intuitive data visualization. 3D charts and lots of busy-ness tend to obscure information rather than make it easier to digest.

Determining whether a metric should be displayed as a bar chart, trend line, spark line, or table should be driven by the nature of the metric and the preferences of the decision-makers who will be using the dashboard. Make no mistake: their preferences matter. Our experience has been that there can be a fair amount of trial and error that takes place before you find the ideal dashboard format for some decision-makers. A helpful shortcut can be to get your hands on dashboards they already use and rely on. Aligning your style to those can streamline the process of acceptance and adoption. Certainly don't assume that they'll accept with open arms any old batch of charts and tables you put in front of them.

TALKING ABOUT RISK

The charts and graphs we generate rarely use multiple colors, and we prefer to reserve red, green, and yellow for those instances where there is a very specific need or purpose. The reason we try to avoid using these colors is because people too easily infer "goodness" or "badness" (or "ho-hum" in the case of yellow) from them. Data should convey meaning on its own. Having said that, if an organization explicitly defines thresholds that delineate levels of concern for a specific metric and want at-a-glance recognition of a particular condition that may need a decision (like in a dashboard), then that is a great time to use these colors.

SUMMARY

We have covered a lot of ground in this chapter and we could have, quite literally, made it a book unto itself (foreshadowing perhaps?). Regardless, even though we haven't gone into great depth on individual metrics or provide as many examples and how-to's as we'd have liked, we hope we have armed you with new insights and clarification regarding information security and risk management metrics. If organizations adopt the GQM framework we've outlined (or some derivative of it) and are thus better able to manage risk more cost-effectively through its metrics, then we'll consider the chapter a success.

There is one last point we would like to make that may seem obvious, but we'll make it anyway. Implementing a metrics program based on what has been described in this chapter might seem overwhelming. That being the case, don't try to eat the elephant all at once. Think about your organization and where it might benefit most, metrics-wise, and start there. Then evolve the program over time as success builds and as you overcome the inevitable hurdles you will encounter. Becoming a metrics-based organization does not happen overnight.