PART I

Introducing Security Metrics

CHAPTER 1 | What Is a Security Metric?

So you are ready to set up a security metrics program—or maybe you're not quite ready, but you're curious about how you can better measure and improve the security of your organization. You may be looking for new ways to show the value of security to senior management. Or perhaps you just want more visibility into security operations. You may be worried about compliance with laws or regulations that require your organization to be more accountable for the specifics of security management. Whatever your reason, you are ready to learn more about how to develop and benefit from IT security metrics. Before you dive into those details, however, you need to understand the role of metrics in the security world.

The past few years have seen increasing buzz around security metrics. Several books as well as numerous industry articles, reports, and white papers have been devoted to the benefits of measuring IT security. Security metrics have become a hot topic so quickly that some might assume we have only just discovered that we can measure what we do. But this is not accurate, of course, and well-known security metrics such as annualized loss expectancy (ALE), total cost of ownership (TCO), and quantitative and qualitative risk assessment have been used by security professionals for years.

What is new in security metrics is the growing understanding that many of our traditional efforts at measurement are unsatisfactory. They do not give us the information we really need to support decisions and articulate the value of security activities. And they are not adequate for the changing security landscape of more subtle threats and increased accountability and scrutiny. The growing consensus is that we must measure better and consider new and innovative ways of analyzing the metrics data we already have. The purpose of this book is to add to the IT security metrics conversation and help you achieve the goal of better measuring and articulating the value of information and IT asset protection.

When I advise clients on how to develop an effective security metrics program, I usually face some immediate challenges, not the least of which is that, although people generally understand metrics, it is often localized to their immediate concerns. We tend to measure only those things that we deal with regularly, and eventually we decide those are the only measurements that matter. For example, every morning I make coffee, carefully measuring several scoops of ground coffee and several cups of water into a French press as part of my daily caffeine ingestion ritual. I care about these measurements because they directly affect my morning. I don't think about how these measurements are related to other metrics, such as the proper acidity and nitrogen levels for growing coffee or the optimal temperatures and durations for roasting it. I depend on others for these measurements (although if they are incompetently performed, I find another source for my coffee).

Metrics, both for coffee and for IT security, involve many local and tactical efforts that become increasingly interdependent and strategic as they begin to affect larger systems. I may not perform measurements outside of my local context, but, if I'm smart, I will try to understand more about them so that I can make the best decisions. And others will do the same. Understanding what makes good coffee beyond just grounds-to-water ratios will help me maximize my consumption experience, while understanding how I measure and drink my morning beverage will help coffee producers show value and compete

for my business. It is no different for IT security. I may not measure security beyond analyzing the contents of my firewall logs, but if I don't understand how others measure security or other business values, I will be less able to use my data to make (or help others make) good decisions. And if I can learn to understand how other stakeholders in my business measure success, I can use my security metrics data to help them be more successful in their operations, thereby demonstrating the value of my own activities.

As security becomes more complex and pervasive, and security professionals are held responsible not only for protecting company assets but also for contributing to its financial and competitive success, information about how IT security operates will be more globally and strategically relevant. As a consultant, I am exposed to a wide variety of requirements and environments that have proven the value of a broad understanding of security measurement. I advise people to take a big-picture approach to security metrics.

To return to my analogy, if your livelihood depends on coffee, you need to understand more than just the mechanics of a cup of joe. Likewise, if you are a chief information security officer (CISO), you need to know more than just how many events the firewall logged yesterday or how much one vendor's firewall might save you over another. Measurement is also about understanding why we want to measure something in the first place, what it is that we actually want to measure, how we can measure it, and what to do with the data we collect. So let's begin by taking a look at metrics and measurement in general.

Metrics and Measurement

You might want to implement a security metrics program for some immediate reasons, including justifying the value of your activities to management or improving your ability to control and secure your infrastructure. But at the heart of your reasons lies the single reason why we measure anything: we want to understand it better. This is a key point that will inform the rest of the book and your efforts to implement metrics within your own security program.

You measure security to understand security. This statement may seem simple, but it is more difficult to put into practice than it seems. I know clients that have established metrics programs and yet still struggle with understanding their security efforts. This often occurs because a client's metrics program is actually a data collection program and not measurement-driven at all. These metrics programs remind me of the giant warehouse in the Indiana Jones movies, where the government stashed away and subsequently forgot about every cool mystical device that Dr. Jones worked so hard to procure. Collecting security data is critical to any effective metrics program, but without a context for that data and an idea of why you collected it and what you intend to do with it, you might find yourself limited to describing your measurement only in terms of terabytes of log data and the shelf volume occupied by auditor reports.

Metrics Are a Result

One of the common mistakes people make when setting up a security metrics program is to focus too much on the metrics themselves. Some of the blame for this focus can be assigned to semantics, because the industry has adopted the term *metrics* in favor of the term *measurement*, which would better describe what we are trying to accomplish. I'm guilty as well, as evidenced by the title of this book, but I try to choose my battles, and a book on IT security *measurement* would be dissonant and perhaps confusing.

The important point to emphasize is that security metrics are a *journey* and not a *destination*. Once you have established a security metrics program, you must ask yourself how the results of the program have improved your understanding of your security systems and processes. Understanding is not diagnostics. Knowing year after year that some percentage of your users' passwords are easily cracked or that the ratio of vulnerable to secure Internet-facing hosts hasn't dropped below 1-in-4 reduces some of the uncertainty regarding your IT security effectiveness, but if the information has not enabled you to improve that effectiveness, something is missing from the program. Even if the security has improved, if that is all you know and you cannot say why the improvement occurred, your metrics are not giving you any more value than if you were struggling over why your security was getting worse. Metrics are conceptual data repositories—they define and standardize information. Metrics do not organize that information into knowledge, any more than well-defined word entries will transform a dictionary into literature. Only people can accomplish these things.

Measurement Is an Activity

The point of security metrics is not to collect a lot of data. A small set of data, understood well and applied regularly, is much more valuable than a mountain of data left untouched on shelves or hard drives and gathering real or virtual dust. The true benefits of metrics come when the data that they represent is the end result of meaningful activities, actions that we take to accomplish a goal or a task. Metrics are, or at least should be, the records of our observations. Measurement is the activity of making observations and collecting data in an effort to gain practical insight into whatever it is that we are attempting to understand. The distinction is important, because metrics bring not just information about IT security, but also costs and risks.

Collecting metrics data for the sake of collecting metrics data is not measurement unless the purpose of the activity is to mine historical data for interesting patterns as a research exercise. I actually love this type of measurement and think it is valuable, but most of the clients I work with that collect security data do not do so for academic reasons, and their security data is rarely analyzed historically or experimentally. More often, the benefit of collecting security data is directly related to the ability to claim that a lot of security data has been collected. Having great amounts of data at hand can be comforting, providing a reassuring sense that we are on top of things even if we have no real clue about what the data reveals. Collecting all this data may even serve as ammunition in support of organizational rivalries as people strive to collect more data than the peers, supervisors, or groups with whom they interact and compete.

The challenge is that security metrics are inherently risky, as is anything that allows you to understand something better than you understood it before. Knowledge may be power, but it often carries with it certain demands and obligations, not the least of which is that you may have to consider new ways of looking at your environment that can be quite uncomfortable (or expensive). In addition to the simple questions of overhead that come with collecting and storing metrics data, there is the implication that whatever security data you collect now constitutes something that you are aware of as an individual or an organization. Collecting data regarding the vulnerabilities in your systems implies that you now know how vulnerable they are, because that information is in a report, either from an automated tool or from a consultant, that may be sitting on the shelf behind your desk or on your hard drive. In the event of a security breach, those metrics may even become discoverable should your organization face litigation.

Whether you read the report or understood it is immaterial: you collected the data and increased your knowledge. Knowing about the problem and not having acted upon it, leading to a security breach, however, could actually end up more damaging than the ignorance that existed before you ever gathered the data. Many security managers don't consider the idea that the data they collect becomes a matter of corporate record and possibly subject to e-discovery. Unused metrics data simply adds insult to injury. You still get hacked, but you also lose the resulting lawsuit because you "knew" you could get hacked based on your security metrics data. This is an important consideration for security metrics that is only beginning to be discussed in our industry.

My point is not that metrics are too risky and that we should strive to know as little as possible about how our security is functioning. It is that if you collect data and do not use it, you do not have a security metrics program. Measurement without analysis and action wastes time and money and contributes to uncertainty and risk rather than reducing them. We need to know more about our security operations. The value that comes from understanding our security processes far outweighs the risks associated with that knowledge. But metrics must be based on a sound strategy for security measurement and applied understanding, and not about hoarding data that we never intend to look at again, much less put to productive use. Instead, security metrics should be seen as part of a business process that continually seeks to improve the protection of enterprise information assets over time.

If you are undertaking a security metrics program, you should do so with the same eye toward risks, costs, and benefits that you would approach any other business process. For every metric your organization collects, you or someone must understand why that data is being collected and what decisions the data will be used to support. And someone should be assessing the costs and benefits of collecting the data. It is fine (and often useful) to collect exploratory data that is not associated with any particular objective, but research metrics should also be understood and should eventually lead to new knowledge and insight for your company.

As you put metrics into place to explore your security operations, ask yourself whether you are prepared to act on the knowledge you gain through your measurement program, even if it is unexpected or imposes new obligations and requirements on your security operations. If you are not ready to act on what you discover, metrics are only going to compound your problems.

Security Metrics Today

Increased interest in IT security metrics notwithstanding, the security industry already uses several commonly recognized metrics. Some of these metrics are cornerstones of security practice for vendors seeking to market their products and for security managers trying to improve security and reduce risk. The problem is that many of these metrics have limitations that make them misleading indicators of security effectiveness.

There are plenty of arguments about what makes a good or a bad metric, and I will explore some of these arguments throughout this book. I believe that any empirical measurement that helps an organization reduce uncertainty is a good metric. I do not believe that a metric should be discounted simply because it is not quantitative or specific, or that a metric is good simply because it is easy and unambiguous. Any measurement becomes problematic when it is conducted poorly and when those measuring are not sufficiently critical of their own methods. Problems that can arise from unsophisticated attempts at measuring security can include issues of data quality, empirical rigor, or the fact that the metrics are used in immature or misleading ways. The following metrics all suffer from one or more of these problems.

Risk

Risk is a foundational concept in IT security. At the heart of any security-related question is the deeper question of what risks we assume by making a certain decision or taking a particular course of action. Of all the phenomena that we care about understanding as security stakeholders, risk would seem to be at the top of the list. But as critical as an understanding of risk is, it is often one of the most poorly understood concepts. Information security practitioners typically use terms such as *risk assessment*, *risk analysis*, and *risk management* as generalities in which the definition of risk is often assumed or taken for granted. In IT security, risk is typically associated with some harm or loss to systems or data, but this definition is too general and not universally accepted or consistently used. Instead, risk is usually bundled into some combination with other generalized issues of threats, vulnerabilities, and parameters that are often equally imprecise until we are left with a fuzzy concept that can change across organizations and implementations. This makes risk difficult to measure consistently in security, and it doesn't help that many vendors confuse the meaning of the term or misuse it when they try to sell their security products and services.

IT security's approach to risk can reflect the relative immaturity of the industry and our responses to the professional challenges we face. Our understanding of risk is something of a catch-all, and we rarely feel the need to be clear about what we actually mean when we discuss it. We use the term *risk* to describe many different phenomena that we know can affect our security, but that we have not yet explored and defined.

When you mention risk in an IT security context, everyone will nod in agreement, but you can never be sure that everyone is thinking of risk in the same way. Risk can, after all, mean a lot of things. Consider a mature industry such as finance and the definitional problem is put into perspective. Ask a finance person about risk, and she may require more clarification about what you actually mean. Are you referring to *endogenous* or *exogenous* risks—risks from events within your control or risks that come

from outside of your control? Or are you talking about *systematic* or *unsystematic* risks—whether or not the risk is subject to chance as defined by some probability curve, or whether the risk is non-probabilistic? These are just a few of the specific characteristics and types of risk that might be referred to in a discussion of formal risk management.

You may find that, in the eyes of your colleague in finance, you need to do a bit more homework before you are ready to consider measuring your risk. Immaturity is a natural thing. Insurance and finance companies didn't always measure risk with the sophistication that they do today (and in the wake of the recent economic crisis, some will argue that they remain immature in some ways). Measurement improves with practice and discipline, and the more security pros actively attempt to measure and understand our operations, the better we will get at our assessments.

Our somewhat naïve definition of risk in the context of IT security is mirrored by the lack of rigor we tend to demonstrate in measuring it. Probably the most common method employed in measuring security risk is to use a variation of the "Likelihood × Severity" matrix shown in Figure 1-1. Some version of this formula and matrix can be found in the

Generic Risk Matrix

		Likelihood of Event		
		High	Medium	Low
Severity of Impact	High	"We're Doomed!"	Bad	Outlier
	Medium	Bad	Not Good	Error
	Low	Annoyance	Typical	"Whatever..."

Figure 1-1. Generalized risk assessment matrix

majority of discussions, books, and training programs regarding IT security risk assessment. The matrix may be more complex and contain different scales, weighting factors, heat map colors, or other bells and whistles, but they all are derived from the same concept. The idea is that you estimate the likelihood that something (usually a technology system) will experience a negative security event, and then you estimate the severity of that event in terms of how badly the system is impacted. The results are used to populate the matrix and give you a prioritized summary of your risk. The matrix is simple and makes intuitive sense, which is likely why it has persisted for so long. Nevertheless, as an instrument for measuring risk, it is pretty limited, certainly too limited to justify the enormous amount of stock that we put into it in support of some of our security decisions.

While it has problems as a measure of actual risk, the matrix can be quite effective as a targeted opinion poll. It allows security subject matter experts to prototype quickly what they believe to be their biggest security problems. You see this type of assessment used all the time in the media, when experts are brought in to clarify and provide opinion on current affairs and events. These individuals have knowledge and experience that should make them more suitable to comment on the topics under consideration than just anyone off the street. Of course, none of this expertise proves that these people are correct, and in fact experts often disagree. The point is that experts should have more informed opinions regarding the areas of their expertise than the rest of us—this is why we have teachers and doctors and attorneys and security specialists in the first place. Their insights can clarify a subject and remove the confusion and noise surrounding it, allowing us to focus on what really matters.

The important point is to recognize that opinion alone can have value, and not to insist that the opinion also represent a fact in order to have merit. A security risk matrix based on expert judgments can be a useful estimate, but it remains a set of opinions about risk. The biggest security problems identified in the matrix are not necessarily the biggest security problems facing the enterprise. The hope is that the true security risks will correlate in some way with the expert opinions of those responsible for security. As I will describe in later chapters, there are ways to calibrate and refine expert judgment to make these opinions less uncertain, but there will always be a margin for error. When we deliberately ignore this uncertainty because we want to pretend we have identified a fact, we lose track of what we are measuring and our matrix becomes misleading and contributes more, not less, uncertainty to our decisions. This result reflects the first of two fundamental limitations involved in this form of risk assessment.

Security Risk Assessments Don't Measure Risk

Consider the standard security risk assessment methodology. Groups of stakeholders are gathered together or surveyed by questionnaire and asked to provide risk scores for probability and severity of occurrence for their systems and data. These individuals dutifully provide the requested data, which is used to populate the matrix. The result is that a measurement has certainly been conducted. We can even claim that the measurement was more or less empirical because it involved observing some phenomena.

The problem is that where we think we measured security risk, we actually measured human judgments about security risk. In more formal measurement terms, we have just developed what is known as a validity problem—what we think we are observing does not accurately reflect what we are actually observing.

Some critics of this simplified form of risk assessment go to the opposite extreme, believing that since you are not actually measuring risk, the entire assessment matrix exercise is worthless. I tend to disagree. Nothing is intrinsically wrong with measuring someone's opinion of something. If such measurement did not produce valuable results, the marketing and advertising industries (not to mention political consulting groups) would have collapsed long ago. The important consideration is that, when the marketing department of your favorite gadget measures consumer opinions on product quality, they do not make the mistake of thinking that they are actually measuring how good the product really is. Security managers could do a lot to improve the quality of their risk assessment activities by simply recognizing this subtle but important point—that they are measuring opinion rather than risk, but that opinion is also valuable. They might then make the risk assessments more rigorous by focusing efforts on improving the judgments that they elicit, perhaps by calibration exercises and the use of confidence intervals, instead of insisting on turning those opinions into hard numbers that look better in a chart.

Measurement Slackers and "Statistical Alchemy"

A second problem with the current state of security risk assessment results from the fact that, whether consciously or not, we all realize that those assessments are a bit lame. Because we realize this, some security practitioners may feel compelled to try to improve on the method, to make it appear more complex or more rigorous than it really is. At their core, matrix-based assessments take two basic parameters—"how likely?" and "how bad?"—and assign three basic levels—low, medium, or high. And these parameters are derived from data sources that are subjective—namely, people. Anyone thinking about the matrix approach in this light realizes that it makes it difficult to approach senior management with "objective" results based on the exercise. But senior management often isn't interested in opinions; they want facts that they can use to make their decisions, and nonfactual results appear to be less valuable.

The security community has two common responses to this perceived limited value of the risk matrix. The first is to label the risk matrix methodology a "qualitative" risk assessment, which, in IT security terms, tends to translate into "Security is fuzzy stuff; you can't really measure it as you do other things, so you can't blame us if our results prove wildly inaccurate." This is, of course, nonsense. It is the slacker way out of the risk-measurement problem, where we manage to justify the use of the methodology while distancing ourselves from any results we might obtain from it. It also gives qualitative research methods a bad name, implying that they cannot be rigorous or empirical, which is also nonsense. This argument actually functions to relieve security managers and risk-assessment team members from having to critique and improve their own measurement activities.

Even worse is a practice I call "statistical alchemy," which involves transmuting one thing into something completely different that is perceived as more valuable. As I noted earlier, the risk matrix generally involves assigning high, medium, and low levels of likelihood and severity to a particular event under consideration. These levels are on what is known as a *nominal* scale. I will address levels of measurement such as nominal, ordinal, and interval measures later, but for now suffice to say that nominal measures function as discrete categories. Hot and cold, good and bad, and high and low are all nominal, meaning that you cannot compare them in terms of value, scale, or ratio to one another. Business decision-makers tend not to like the inputs to those decisions expressed so categorically; they want to see numbers, to know *how much* hotter or colder, better or worse, or higher or lower something is. Numbers add a sense of certainty and importance to observations, whether or not they actually provide those things. Luckily, when a risk analysis is conducted for someone who is expecting to base decisions on numbers, a simple solution is at hand: Just change all the levels to numbers! Now a high likelihood is a 3, a medium likelihood is a 2, and a low likelihood is a 1. The same goes for high, medium, and low severity. This lets you successfully transform statistical lead (an ordinal measurement) into something that may not be gold, but is closer than you were before. Calculating the average of high and medium (medium-high?) is meaningless, but calculating the average of 3 and 2 is not (it's an unambiguous 2.5).

Most assessments that adopt simple numerical categories would not be portrayed as quantitative. Security folks are smart people, and we would see through such a ploy. But more "sophisticated" risk analysis matrices up the ante. Instead of numbers corresponding to high, medium, or low, perhaps they require the specification of a dollar loss, such as "below $25,000" or "above $500,000" in the severity columns. Likelihood levels may be replaced with probability scores, such as "90% likely" or a "0.25 probability" that an event will occur. Additional columns can be included to simulate numerical weights based on the system's environment or the ratio of system functionality that may be lost. Now the matrix becomes something more like a spreadsheet, with the highest risks expressed in estimates of financial loss. It is our same humble risk matrix now dressed up, Pygmalion-like, as something more than it is. And even if those who conducted the assessment are still reminding everyone that the matrix is qualitative, reflecting human opinions and not real numbers, no one is really listening anymore.

So Why Even Use the Risk Matrix?

The real tragedy of the security risk matrix is not that it is a bad method of measurement, but that it is bad to pretend that the matrix measures actual risk. Unfortunately, most users of the matrix in IT security do not give much thought to the importance of that nuance, and they use the matrix to make "risk-based" decisions. Even considering the hedgers who caveat the matrix with the word "qualitative" (and then often go on to treat the results as factual), the risk matrix has become the engine behind some of the most common security–risk-assessment methodologies today.

It seems that new variations of the matrix are developed every year at significant effort and cost. Often these methodologies are used as the organization's formal risk assessment and management methodology, as required by some compliance frameworks. In these cases, the matrix does not act as an initial prototype of risk measurement that leads to more questions and metrics, but as the end result of the risk assessment process. It is as if an insurance company made underwriting decisions based on the experiences and opinions of a team of actuaries and never bothered to verify whether those opinions were correct before handing out policies. I don't advocate abandoning risk matrices as a means to support security decisions, but I do think that these tools should be used for at least two different purposes than they are used today.

Assessment Prototyping A security risk matrix is, as I mentioned, a good barometer of people's thoughts and perceptions regarding risk. And since the methodology expects you to ask risk questions of people who are responsible for the systems under review, knowing what these experts think about the risk levels of the systems they manage can be valuable data.

Some of the best value comes when we use the matrix as a means of prototyping further risk assessments. Too often I see organizations that have undertaken a general risk assessment methodology and accept the results without ever asking the all-important question "Why?" Why is this system so likely to be compromised, and why is the impact so severe compared to the other systems? Instead of simply accepting the rating, asking why encourages security managers to think about follow-up questions, which lead to more measurements. Asking these questions does not mean you disagree with or challenge the risk rating, but that you need to understand why the claim was made so that you can effectively respond to it. As the first step in defining the data we need, the tests we must run, and the metrics we must define to assess our risk, a risk matrix can function quite effectively and not be ruined by expectations that should never have been laid upon it in the first place.

Measuring Differences in Agreement Another great use for a risk assessment matrix is to compare what different people in the organization think about risk. Rather than treating the matrix as a reflection of reality, the scores used to populate the data can be used to identify areas where everyone is in agreement or everyone varies widely in the opinions that they hold. This, too, can provide valuable data, particularly if major disagreements exist over the importance of particular systems or how much the organization would be hurt should they be compromised.

This approach encourages the assessment team to expand the pool of experts from which they collect data. You might find, for instance, that the e-mail administrator is far more concerned with a loss of service to users' inboxes and rates e-mail storage as a relatively low risk, but the compliance officer responsible for records retention and e-discovery is far more concerned with compromises in the e-mail archiving system. As with prototyping, this use of the risk matrix serves primarily as a means to discover where the organization should concentrate its risk assessment efforts, including where to conduct more sophisticated and robust measurement activities.

Security Vulnerability and Incident Statistics

Measure for measure, the data most often collected for the purpose of understanding IT security involves system vulnerabilities and efforts to compromise them. System vulnerability statistics are produced when an organization runs a security scanner on its network, when new exploits are identified and released to vendors and the public, and when organizations release reports resulting from industry surveys they have conducted or analyses of security data they have collected. Incident statistics come from system logs, intrusion detection and prevention systems, and industry surveys and analyses. These numbers are often used as general indicators of the current stat of IT security.

A Parade of Horribles

I recently read a vendor-sponsored industry research report on Internet security trends. The report included a scatter plot chart that showed the number of reported product security vulnerabilities over time. It showed an obvious positive correlation as the number of vulnerabilities increased steadily over the timeline of the graph. The report concluded that Internet security was getting worse (a trend that certainly justified the sponsorship of the security vendors who subsidized the research study). The problem here is that measuring Internet security by the number of reported vulnerabilities each year is like measuring male virility by the number of prescriptions written to treat erectile dysfunction. If I charted these prescriptions on the same chart as security vulnerabilities, it would appear that male reproductive capabilities were in rapid decline during the last decade or so and that the human race might be in trouble. Both analyses ignore more data than they include. From a security perspective, the mere addition of hundreds of new technology products every year could be enough to account for the increase in reported vulnerabilities.

Counting and analyzing technical vulnerabilities and the attempts to exploit them are important aspects of any IT security program. But if you make security vulnerabilities the primary data you use to measure your security, you cannot help but distort and skew the results. Relying too much on vulnerability data contributes to fear, uncertainty, and doubt (FUD) rather than rational attempts to analyze and improve security business processes. When that analysis is also sloppy, as in the security report I found, the problem is compounded.

A Thousand Walled Gardens

Vulnerability and incident data reporting is not problematic only because of its tendency towards hyperbole. As a measurement, it is inconsistent because it occurs in too many places and in too many ways, without sufficient aggregation or normalization of the data. A company running a vulnerability scanner against itself is not likely to share the information it gathers with other companies or even with other groups inside the company. Vendors and consultants publishing this information for a fee or as a way to promote their products and services are unlikely to be forthcoming, because the data represents valuable intellectual property. This reluctance to share data and the lack of

effective systems to facilitate sharing among organizations make it that much more difficult for academic researchers and public institutions that might want to distribute the information. The result is that most organizations have no data to rely on other than what they collect and no real way to compare their data with anyone else's data.

The most common question I am asked by clients from a security perspective is how well they stack up compared to their competitors and other companies; I am always forced to admit that I cannot provide a satisfactory answer. Of course, there have been efforts to share security data, with efforts ranging from high-level surveys and studies such as the Computer Security Institute's annual *CSI Computer Crime and Security Survey* and a host of studies by vendors and market analysis firms. Other technical efforts have attempted to normalize vulnerability data, including the Common Vulnerabilities and Exposures (CVE) dictionary and the Common Vulnerability Scoring System (CVSS). But while these resources help with general understanding, they do not reflect anything close to the common metrics and shared data that exist in more mature industries such as insurance, transportation, or manufacturing.

Annualized Loss Expectancy

If vulnerability-related statistics are among the most commonly collected measurement data in security, ALE is the most commonly used conceptual metric. ALE refers to how much you think you will lose as a result of security incidents. Where risk assessment matrices are used to compare and prioritize risks qualitatively into cells in a table, better to identify where to focus security efforts, ALE is pitched as a fully quantitative metric, complete with formulas and other statistical goodness.

The formula is expressed as ALE = ARO × SLE, where ARO is the annualized rate of occurrence (how often you expect to experience the loss in a given year) and SLE is single loss expectancy (how much you expect one incident of the loss to cost you). Suppose, for example, that you have a server worth $10,000 (system and data combined) and you estimate a 25 percent chance that the server will be successfully compromised as the result of a zero-day exploit in the coming year (ARO = 0.25). Each time the server is compromised, you estimate that you will lose $5000 due to remediation costs and the exposure of the data stored there (SLE = $5000). Your expected annual loss is then ALE = 0.25 × $5000, or $1250 each year. Theoretically, you have now identified your security budget for that particular server, as you should not spend more protecting the asset than you would lose should it be compromised.

I find the ALE formula interesting because it is unique to the security industry. You might expect that it was borrowed from the insurance industry, which has a much longer history of risk assessment and management. In fact, as far as I can tell, the metric first emerged in the 1970s as part of Federal Information Processing Standards Publications (FIPS PUBS) published by the National Institute of Standards and Technology (NIST). And in those three decades, the metric and the way it is used have hardly changed, while ALE has developed into perhaps the most common single measurement in IT security. Unfortunately for security managers, ALE is a poor metric.

Expectations vs. Probabilities

I am certainly not the first to critique ALE as a security metric, and it surprises me how the formula continues to gain and maintain acceptance as an IT security standard by professionals who should know better. Like general matrix-based risk assessments, ALE relies on data that is often completely fabricated. This is reflected in its name, which implies human expectations. If it were called Annual Loss Probability, the formula would at least imply that the results were based on more concrete data. Like the risk matrix, ALE measures what people *think* rather than objective reality. The people in question may know a lot about the systems they are asked to review, but when a risk assessment team polls its members to populate the ALE formula, they are soliciting opinions. ALE is a perfect example of statistical alchemy. Unlike the risk matrix, which, though flawed, presents data in a categorical context that does not necessarily imply how things will actually turn out, ALE pretends to show you probable outcomes.

ALE deals in opinions and expectations primarily because IT security does not have the data necessary to define actual probabilities. The discussion of security vulnerability and incident data showed some of the weaknesses involved in collecting meaningful security data. Part of the problem is that most organizations do not have systematic programs for collecting and analyzing historical data even for vulnerability and incident data, much less the impacts and losses that they have experienced as a result of security breaches. In many cases, organizations are not even able to detect or track events that would lead to this data in real time. In those rare cases in which an organization is detecting, collecting, and analyzing this data, there is no collective industry mechanism by which this data can be shared, even assuming that the organization wants to share it. Most do not. Industries such as insurance function because they have made a science of collecting and sharing data regarding the risks that the industry faces as a whole. IT security has not matured to a level at which we are able to do this—one of the many reasons that real, verifiable security metrics are becoming more important to everyone.

What Have We Got to Lose?

The other big problem with ALE is our lack of understanding about what constitutes loss. ALE can function only by assigning dollar costs to events. Therefore, the metric tends to focus on those scenarios in which the system in question is rendered inoperable for some period of time, where time must be spent to clean or repair the system, or when the value of the data residing on the system is negatively impacted through theft or exposure. (Assigning value to our data is a completely different problem that also complicates our ALE results.)

ALE does a poor job of estimating the risks associated with intangible losses to such things as brand or reputation. The model is blunt and inaccurate, and the moment you try to add nuance or sophistication to your analysis, it tends to break down. Part of the problem is a lack of awareness of our security environments. Just as organizations have a hard time gathering data on attacks and events, they often do not have a sophisticated awareness of what losses they might incur.

ALE tends to focus specifically on technology systems, because they are the easiest to model. We mislead ourselves into thinking we can understand our losses based on an analysis of hardware, software, and data because we can calculate their value even if that means only factoring in how much we paid for them. But this valuation is often the least useful for risk assessment, because what we really want to know is not direct replacement cost, but rather how the loss of an asset creates other losses such as those involving productivity, efficiency, or competitiveness. Determining these losses brings us right back to our limitations of data and awareness and forces us to rely not on verifiable data and probabilities but on more or less educated guesswork.

Return on Investment

Return on investment (ROI) is a security metric that has to do with calculating how much benefit (usually described in financial terms) will be gained from an investment. IT security borrowed ROI directly from the business world, where the idea of taking more out of your efforts than you put into them (also known as profit) is of central importance.

From a security perspective, we usually refer to ROI in a couple of ways. First it is related to ALE, which defines the expected security losses incurred in the absence of any preventative action. If an organization takes preventative action, the relationship of the cost of the action to the expected losses defines ROI. If, for example, you expect to lose $10,000 in a security incident and prevent that loss by spending $1000, your ROI is $9000. If you spend $20,000 to prevent the same event, you have a negative return of $10,000. You can beef up the measurement in other, fancier ways, such as weighting or discounting the return over time, but these are the basics.

The second way ROI gets used is by security vendors as a means of marketing products. The vendor builds models to show how an organization that buys its product will end up getting a great ROI. The vendor may include ALE analyses that show how the product reduces loss as well as ways that the customer can benefit from improvements in efficiency or productivity. The vendor can then use these ROI figures in conjunction with pricing and support options to show a customer that the product provides the most bang for the buck.

ROI in an IT security context also qualifies as statistical alchemy, because it misleadingly tries to equate different concepts quantitatively. In the finance industry, for instance, ROI might be reflected in the rate of return on a monetary investment in which a borrower agrees to pay a lender for the use of the lender's money. In capital expenditures in industry, on the other hand, the ROI has to do with profit, the amount of additional money that can be made through the use of a fixed cost asset over time. Security does not really function in either of these ways, because security activities are not undertaken as a profit center (unless they are provided for somebody else as a business). IT security has to do with loss prevention, much like physical security mechanisms such as locks, fences, and guards.

The reason IT security is portrayed as an investment has to do with marketing. The main use of security ROI figures is to convince someone with money to give that money to someone else, and most people feel more comfortable about giving away money if

they think they are investing it. This is why security ROI is used a lot by security managers who make business cases and by security vendors that sell products—both have a vested interest in convincing someone to give them money.

As with the security metrics discussed previously, the biggest problem with ROI is the data that goes into the equation. If the data is unreliable, the metric is equally meaningless. ROI has an additional stigma in that, because it is used directly to influence financial decisions, it encourages people to manipulate the data to achieve the outcome most favorable to them. This makes ROI doubly unreliable, because you not only have to account for incomplete and subjective data, but now you must consider whether the metric is not just inaccurate but deliberately misleading.

Total Cost of Ownership

Where ALE attempts to measure losses associated with IT systems and ROI attempts to measure the "profit" derived from them, TCO seeks to quantify the money that must be spent on the system throughout the entire ownership lifecycle, from initial purchase to final disposal.

TCO was first developed by the Gartner Group in the late 1980s as a way of helping its analysis clients compare vendor products. TCO is designed to take a more holistic view of the cost of a particular system and to include factors that may not be reflected in the purchase price, including the following:

- Central system components such as hardware and software
- License and support fees
- Supporting infrastructure (space, power, environmental controls)
- Installation and maintenance
- Training and expertise
- Security and audit
- Hidden costs

TCO in IT security is designed to mirror TCO in other industries. For instance, most of us realize when we buy a new car that we have to factor in long-term costs such as insurance, maintenance, and fuel. Security TCO attempts to make similar costs associated with data protection systems more visible, so that a picture of the actual costs of a system is revealed.

TCO is more likely to bring some quantitative rigor to the results of the metric than ALE and ROI, because some of the parameters of ownership have more supporting data. But this strength also limits the utility of TCO as a broad security metric, because it applies only to security purchases and not to the measurement of the IT security process. TCO can help you to understand how much a security product will cost over its lifetime, but that doesn't tell you whether or not it will meet your security needs.

Security TCO cannot escape the data uncertainties of other common metrics. Since the security world can't agree on how to track or measure the impact of security incidents, many costs remain hidden and unavailable for inclusion in the analysis. TCO, like ROI, has also been co-opted by security vendors that recognize it as a purchasing decision support metric. These vendors spend a lot of time developing TCO statistics to influence CISO purchasing decisions directly as well as to gain CISO buy-in and support for larger infrastructure purchases. As much as TCO can be a tool to help customers compare solutions, it is also a primary means by which vendors compete with other vendors. No vendor is going to claim higher TCO than a rival when chasing the deal and the motivation for manipulating data and conclusions is high.

TCO can be a useful comparative metric. When factored with other measures, it can support some specific security decisions, including larger IT infrastructure purchases where the vendor has had the foresight to include TCO measures from a security perspective. But TCO does not measure security operations, and the fact that it is one of the most common metrics used in the industry speaks to how much we can improve on our current state.

The Dissatisfying State of Security Metrics: Lessons from Other Industries

The limited number of metrics commonly employed in IT security and the limitations presented by the metrics themselves mean that we do not have the appropriate tools to understand or improve our security systems. This bothers me, because there is no reason that we should not be doing better. We are an industry full of smart people who care deeply about protecting our systems and data. We should be able to measure the results of what we do every day more effectively. Security is not the first industry to deal with complexity, uncertainty, or risk, and if you are considering setting up your own security metrics program, it pays to understand how other professions have dealt with challenges similar to our own and how to overcome the shortcomings in our own efforts.

Insurance

The insurance industry has been professionally managing risk for several centuries, and the security industry could learn a lot by taking cues from its older and wiser forebear. The single most important asset in insurance is data. Data allows insurers to understand the probabilities involved in events against which their customers seek to be protected.

Data collection for insurance purposes dates back to the seventeenth century, when information about everything from mortality rates to shipping routes began to be collected and traded, often in London coffeehouses such as that of Edward Lloyd, of Lloyds of London fame. The data that was collected was subjected to

relatively new and innovative statistical analyses that allowed insurers to predict the likelihood of loss and thus set policies and insurance rates accordingly. Today insurers can issue policies for just about everything from your car to specific body parts, adjusting rates accordingly based upon probabilities gathered from observations of all aspects of life.

Security managers can find it challenging even to provide current and accurate configuration data for the systems they operate. Without data, you cannot even describe daily security activities, much less generalize to how your security functions across the company or across your industry. It is unsurprising that when I first entered the IT security industry, I heard a lot of talk about insuring security risks, and now, ten years later, we still have not been able to make it happen. The insurance industry provides us with the first lesson of IT security metrics:

Security Metrics Lesson #1 Your security metrics and your subsequent risk-management decisions will improve as you improve your capability to collect, analyze, and understand data regarding your security operations.

Manufacturing

The manufacturing industry depends on processes designed to create similar products on a mass scale. Variation in these products is highly undesirable, because it introduces problems of quality, efficiency, and reliability in that which is produced.

Whether the manufacturing process is the injection molding of plastic drinking cups or the assembly line activities of an automobile plant, manufacturing industries must ensure that each product is free of defects within strict and predefined parameters. At the same time, the manufacturing process must be constantly monitored and improvements made to the efficiency and productivity of operations if the manufacturer hopes to compete with other manufacturers.

The manufacturing industry has been studying how to improve its processes for nearly as long as the insurance industry has been managing risk—at least as far back as the famous economist Adam Smith's description of the benefits of division of labor in the typical English pin factory. In the early twentieth century through the end of World War II, process experts began applying sophisticated statistical models to the manufacturing process in an effort to increase efficiency and quality in the products created. In the decades since, manufacturers have conducted much research into quality management and statistical process control methods that allow for high degrees of consistency and standardization even in highly complex production systems such as microelectronics and biotechnology.

Your security program may not function exactly like an assembly line, but unless your security operations are very different from most others, you are also not treating your security as a true business process. You may have security processes in place, but it is unlikely that these processes have been formally deconstructed, mapped, or analyzed at levels of detail sufficient to implement statistical controls on the activities involved. So it is likely that many of your security activities remain somewhat opaque

and unclear even within your own organization. You can and should consider many techniques and methods from the process control research literature to understand and improve these processes. Along with the need to collect more data, a process approach to security is the most important improvement strategy that you can undertake, and this is the goal of the Security Process Management Framework described later in the book. For now, we can take from the manufacturing industry our second lesson in security metrics:

Security Metrics Lesson #2 Security is a business process. If you are not measuring and controlling the process, you are not measuring and controlling security.

Design

I am a social scientist by training, so I sometimes find myself at odds with other security metrics advocates who believe that only "hard facts" expressible as numbers should be counted as effective metrics. It often seems to me that one of the end goals of IT security is to rid ourselves of the "problem" of human behavior—if we could just automate everything, users would have no choice but to behave properly. In academia, this is sometimes referred to as "technological determinism" and reflects a state of affairs in which technology rather than people is the primary driver of human society.

No one understands how misleading this view of the world is better than technology designers who deal every day with the consequences of not understanding how central people are the development and use of that technology. What this means for security metrics is that if you are not making an attempt to understand the social, organizational, and even cultural aspects of your security program, you are missing at least half of the picture.

When qualitative measurement is brought up in the context of security, it is often a euphemism for data that is conceptually too "soft" and unscientific or logistically too difficult to collect to be useful. This represents a gross misunderstanding of the purpose and methods behind the science of qualitative inquiry. Designers rely on a variety of "soft" research methods in their work that would likely make believers in hardcore quantitative security metrics cringe or at least roll their eyes disapprovingly. Designers may talk in terms of context, social norms, and even empathy as part of their measurement process, which they are more likely to refer to as research than measurement. (I'll cover the distinction as it pertains to security in the next chapter.)

Design researchers and the companies that employ them use a variety of rigorous qualitative methods such as survey research, ethnography, and narrative analysis to gain insights into areas of human behavior that simply cannot be analyzed any other way. These researchers study everything from people's shaving habits (to create better razors), to the way people use their kitchens (to create better smart appliances).

In security, we often go the opposite way, studying the technologies in an attempt to create better human behaviors. But IT security is inherently a social and organizational phenomenon that involves the use and misuse of technology by people who are not so easy to understand or control. Understanding does not come from deliberately

ignoring what you need to understand, because those things are perceived to be too difficult or expensive to measure or because they do not involve something you can easily count.

The canonical example is social engineering, which has been a bane of security managers since before IT security had to worry about the problem. Whether the deception comes through human interaction or technical hybrids such as phishing attacks, trust trumps technology every time. I find it sadly amusing that our industry recognizes the threat of social engineering but, beyond lip service to training and policy, the main response is often to go right back to an attempt at a technical solution to the problem. So I'll offer a third lesson as you consider your security metrics strategy:

Security Metrics Lesson #3 Security is the result of human activity. Effective measurement programs attempt to understand people as well as technology.

Reassessing Our Ideas About Security Metrics

The security metrics we use today are insufficient to carry us forward into the future of our profession. Security practitioners must develop more sophisticated approaches to security processes in general and measuring and assessing those processes in particular. The experiences of just a sampling of other industries hold valuable lessons regarding how we should think about data, process, and people in approaching our next-generation security metrics. As you develop your own metrics programs, you can and should apply these lessons in several ways to maximize your success.

Thinking Locally

Although it is true that the security industry as a whole is going to need to pull together on metrics, particularly in the areas of common measurement and performance indicators and better sharing of data regarding security operations and incidents, most security managers do not have the luxury of becoming activists. As you develop your own metrics program, you should do so with a keen eye toward your local environment, your organization's specific needs, and the resources that you can bring to bear on your measurement activities.

Metrics programs are not required to be large or comprehensive to be successful; they do need to be better than what was in place before the program. If your organization has no security metrics to speak of, you are in luck, because literally anything that you do will improve the understanding of your security processes. One focused metric, properly analyzed and presented, can be the catalyst for a complete change in the way your organization manages its security. So whether your security program is a tightly run ship or an unorganized mess, it doesn't matter. Metrics can help make it better. You won't accomplish everything overnight, but over the course of this book I will try to help you identify measurement activities that are appropriate to your unique situation and environment and that offer immediate benefits.

Thinking Analytically

Much of this chapter has covered issues of data in regard to metrics: the need for it, where you get it, and why data quality matters. But a security metrics program that collects a lot of data without giving much consideration of what it will do with that data is going to fail. When you build a security metrics program, remember that what you are actually setting up is a program for analyzing the data from your security measurements. If metrics were the end goal, many security organizations would already be finished, and not left wondering why all the data they are collecting is having little real impact on their security. Security metrics analysis means identifying tools and techniques that you can use to create actionable intelligence and organizational learning. Analysis and the sharing of your results widely among the stakeholders that you support becomes the key to transforming your security program from a paradigm of static audit and reactive remediation to one of continuous improvement and innovation.

Thinking Ahead

The thing about measurement is that once you begin, it becomes difficult to stop. Metrics lead to knowledge and insight, which in turn gives you ideas about what else you might be measuring. As your initial metrics efforts gel into a formal process and that process becomes an ongoing program, you should be mindful of what you are hoping to accomplish at the next stage of the game.

As we begin exploring some basic techniques for developing metrics and then more sophisticated tools and methods for analyzing the data that you get, start thinking about what you want to know about your security. Chances are, there's a metric for that. But you may not be able to get to all your security metrics goals immediately. The goal is to stay focused on results. You don't want to drown in a sea of metrics that overloads your ability to analyze the data you have gathered, but you want to begin addressing immediate security measurement goals.

The next chapter offers advice on selecting new security metrics to supplement or replace the traditional and less-satisfying metrics described in this chapter. It provides a methodology for ensuring that your metrics stay focused and aligned to your strategic security and business goals.

Summary

As you consider developing an IT security metrics program, remember that metrics are the result of a measurement process built on human and organizational activities and are not an end in and of themselves. Collecting large amounts of metrics-related data without a cogent plan of analysis and alignment to well-formulated goals is ineffective and can even prove dangerous to the organization, because the argument can be made that any data the organization collects regarding security problems implies awareness of those problems and a responsibility to address them. Your security metrics program

should therefore be designed to provide a manageable amount of usable data that your organization is committed to managing and acting upon, including exploratory or experimental research on data collected without explicit purpose.

The security industry already uses several commonly recognized metrics today to measure aspects of organizational IT security:

- Risk matrices
- Security vulnerability and incident statistics
- Annual loss expectancy (ALE)
- Return on investment (ROI)
- Total cost of ownership (TCO)

Although these metrics are widely accepted, they can be severely limited in terms of the value they bring to a security program. Too often, the measures themselves are poorly understood, measuring aspects of security that are quite different from what their users believe they are measuring. Because of the lack of industry-wide information on security practices and incidents, most of these metrics begin with unreliable data that must be supplemented with non-empirical data such as the opinions of specialists. Although this does not mean that the conclusions drawn from this data are false, it does mean that those conclusions must be subjected to more questioning and skepticism than is typically afforded. In some cases, these metrics are abused by those who manipulate the data to provide results more favorable to their individual or organizational goals.

Other industries have faced the same challenges that the security industry now faces in terms of measuring what they do. As you begin your security metrics program, consider the lessons that can be learned from such industries as insurance, manufacturing, and design. The importance of quality data, the focus on security as a business process, and a greater respect for the role of people and social interactions in the security process are all important elements of a successful security metrics program.

Further Reading

Bernstein, P. *Against the Gods: The Remarkable Story of Risk*. Wiley, 1996.

Condamin, L., et al. *Risk Quantification: Management, Diagnosis, and Hedging*. Wiley, 2006.

Fasser, Y., and D. Brettner. *Management for Quality in High-Technology Enterprises*. Wiley, 2002.

Merholz, P., et al. *Subject to Change: Creating Great Products & Services for an Uncertain World*. O'Reilly, 2008.

Taylor, D. "Your Security Log Files Are a Discoverable Liability." www.thecomplianceauthority.com/security-log-files-are-discoverable-liability.php