

The Case for Risk-Based Security

Jack Jones, CISSP, CISM, CISA

October 2007



In the May, 2006 issue of *The ISSA Journal*, the article “*Making the Case for Replacing Risk-Based Security*” made several excellent observations about the challenges of trying to manage information risk. Some of the points I agree on include:

- There’s a severe lack of credible data regarding information security,
- The risk landscape is so dynamic that estimates of risk probabilities could turn out to be inaccurate as time passes, and
- Much of what’s been described by our profession as “risk analysis” or “risk management” has been unsuccessful at consistently getting management support

On the other hand, the conclusion drawn within the article -- that a risk-based approach to security is fatally flawed and, therefore, should be replaced with standards and laws -- overlooks some important considerations.

This article will cast a different light on the issues raised in the May article, and will draw different conclusions -- even from the points we agree on.

Risk Management “Failure”

I agree with the May article’s claim that our profession has struggled with using risk management practices effectively in the past. I would argue, however, that it isn’t the concept that’s failed us, it’s been our understanding of risk and our approach to risk management that’s failed.

If history has anything to teach us, it’s that simply because something hasn’t been successfully accomplished or accurately understood in the past doesn’t mean that it can’t be accomplished or figured out. Nor does it mean that we should quit trying. Speaking from personal experience as well as from talking to numerous colleagues, management support for security dramatically improves as our risk analysis methods improve and as we’ve begun to treat information security risk as a business risk issue.

Business View of Risk

The May article correctly describes risk as having two primary components -- frequency and impact (which, by the way, is true for all types of risk -- investment, market, security, etc.). It goes on to claim, however, that security risk is somehow different from other forms of business risk, presumably because there’s some profit potential in other business risk decisions. That’s not accurate, though. Every decision we make as thinking creatures has a plus/minus component -- a pro and a con. In our world of information risk, the pro position may be an organization’s ability to function in a certain manner in the pursuit of its goals -- for example, the decision to use the Internet as a sales channel.

As for deciding what the “right” amount of security is, it’s naive (or arrogant) to believe that I -- as an information security professional -- am in a position to understand the incredible mix of business issues that determine the right risk-balance for an organization. Running a business requires weighing the various risk types management faces (investment, insurance, product, market, security, etc.) as well as resource limitations and complex value propositions, and then making decisions about where to place their attention and resources. Even though it’s imperative that information security professionals seek to understand the business side of the equation, we are not going to have the breadth and depth of vision into the organization’s unique mix of business issues that executive management has. Combine that with the fact that it isn’t our risk

tolerance that matters, and it should be crystal clear that complaints of being “underfunded” have to be cast in the light of “Compared to what?”. Compared to what we think it ought to be?

Of course, I struggled to get management support for many years. I tried leveraging fear, uncertainty, and doubt. I also tried the old “You have to do it because it’s best practice” card. And although both of these can work for awhile, at the end of the day you come off as being “chicken little” and/or being viewed as lacking perspective about the nature of running a business. I’ve since come to the conclusion that if I believe I’m underfunded, then it’s likely that I haven’t done a good job of communicating risk to the business, I don’t sufficiently understand the risk tolerance of the organization’s leadership, and/or I don’t understand the mix of competing business risk issues, resource limitations, or objectives. It’s my responsibility to see that I’m not underfunded by providing the best possible risk information to management. If I do that, then I can expect to receive an appropriate level of funding given the other business considerations management faces. The funding may be less than I’d like given my risk tolerance, but that’s a personal problem.

Unknowns

Another claim in the May article is that risk can’t be measured because of unknown variables and an absence of data. Interestingly, several years ago I expressed the same concern to two Senior Vice Presidents of the actuarial departments within a Fortune 100 insurance company. They pointed out two things:

- If you’re attempting to understand, assess, or measure a complex subject (such as investments, insurance, security, etc.) there will always be unknowns, and
- All the data in the world won’t enable you to predict the future with absolute certainty. Case in point: they’ve collected terabytes of data over decades and still can’t tell me whether I’m going to live to be one hundred or die tomorrow. They can only cite odds -- odds which, by the way, are subject to known and unknown variables. The future is and always will be uncertain.

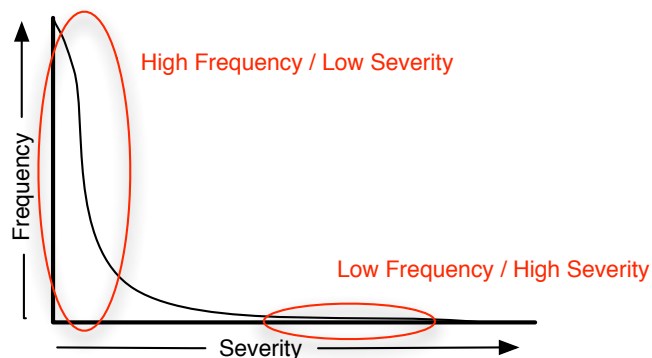
These experts in quantitative analysis and risk also pointed out that no analysis methodology in the world is devoid of human subjectivity. If a human is involved, subjectivity is involved. They stated that one of the key advantages to a structured analysis framework is that it drives greater objectivity into the results.

Clearly, “hard sciences” such as physics have an advantage in using data to “predict” outcomes, but even the hard sciences are limited in their precision by limitations in our understanding of the physical world and our ability to measure it -- and predictions are forecasts, not certainties. There are rare, unpredictable events (“tail events”) in every complex environment due to variables that aren’t known or can’t be predicted.

Risk analysis isn’t a hard science and likely never will be. Consequently, we have to recognize and accept that measurements, estimates, and assessments in our problem space will only lead to better informed decisions -- not decisions without uncertainty. In my experience, business leaders understand and accept this because their other risk domains aren’t hard sciences and they have no preconceived notion that risk analyses (of any sort) are anything but reasoned estimates based on what’s known or suspected.

Rare Events?

Another point worthy of discussion is what the May article referred to as “rare events”. It claimed that rare events (what it referred to as “risks”) are somehow fundamentally different than more frequent events (what it referred to as “certainties”). As already discussed (and as the May article also stated) all events can be described along two continuums -- frequency and severity (illustrated below).



As a general rule, there are far more small-severity events than large-severity events, with some portion falling in between. Regardless of the exact distribution, if we try to categorize events as simply being binary (high or low), then we ignore those that fall in the middle and we inaccurately characterize the nature of our problem. Furthermore, if we “cap” the frequency scale at “1” (i.e., - what the May article refers to as a certainty), then we’re unable to effectively differentiate between events that are expected to happen a little more frequently than once per year and those that are expected to happen many times per year.

A couple of additional things to consider when dealing with the probable frequency of an event:

1. When referring to frequency, probabilities can be expressed as numbers greater than 100% -- i.e., it could happen more than once in the reference timeframe, and
2. If you’re going to measure frequency, you have to do so within the context of a reference timeframe -- e.g., a 30% chance of the event within the next year/day/whatever, or a 200% chance of the event in the given timeframe (i.e., twice in that timeframe). In the absence of a timeframe, a reference to frequency or likelihood is meaningless, as virtually any event is possible given enough time.

Borrowing from the May article, most organizations experience events such as virus attacks on a far more frequent basis than something more significant (on an individual event basis) such as the theft of trade secrets. Both types of events, however, can be expressed as probabilities within a given timeframe. The virus event may happen many times per year (probable frequency greater than once in that timeframe -- i.e., it’s more “certain”), while the theft of trade secrets may typically happen once in a great while (probable frequency less than once in that same timeframe). This approach allows us to more accurately represent the fact that events occur as a distribution along a continuum rather than as some distinctly separate and different category of thing.

Complexity

There’s also concern that the nature of our problem space is too complex to support a well-reasoned analysis. Yes, our landscape is incredibly complex, and no, we’ll never be able to model

it precisely. This doesn't, however, alter the fact that we have to make decisions on a daily basis about how to prioritize our issues, choose between solutions, and defend our conclusions to management. No CISO I'm familiar with has an unlimited budget to get everything done at once, and best practices and standards can't tell me which of the unique combination of risk issues within my organization is most significant. Choosing implies comparison, and comparison implies measurement. And since we're going to have to make choices (and then defend those choices to management, auditors, etc.), we'd do well to use a well-structured risk-based approach.

Absent Data

Perhaps the concern most often cited about security-related risk analysis is the absence of statistically valid data. And there's no question – to-date we've lacked good data regarding many of the scenarios we deal with. This fact, however, doesn't give us a free pass to ignore the probabilistic nature of our problem space and assume “worst case”, or blindly follow someone else's idea of the “right amount of security.” We're forced to make decisions every day -- with or without “hard” data -- so the use of a logical and rigorous analysis method is a sensible means of coming to better-informed conclusions.

The absence of data also begs the question of how “best practice” developers (whoever they may be) determine what to protect against and how much protection to apply. The only rational answer is that they're either operating from intuition or are performing some sort of analysis -- but based on what data and using what analysis approach? If there are infosec gurus out there who hold a secret store of statistically valid data and/or who have some secret process for deriving practices that are ideal for all environments, I'd love to meet and learn from them. I'm guessing, however, that best practices generally come into existence through a combination of theory and informal experiential data -- sometimes with a hint of vendor agenda thrown in for good measure.

Let's face it – any time we estimate the risk associated with some scenario, we're drawing a conclusion by making a series of estimates related to asset value and liability, threat level, and control state. This is true whether we're operating from intuition or from a formal methodology, and whether we're articulating our conclusions using subjective statements or quantitative numbers. The problem is, without a structured approach to analysis we introduce a greater probability of error and inconsistency. Risk analysis methodologies exist primarily to facilitate consistent, rigorous consideration of the information we have available to us or that we can reasonably estimate. They also help to drive bias, myth, and omissions out of our conclusions and recommendations.

Probabilities vs. Prediction

*“Prediction is very difficult, especially about the future.”
(Nobel Laureate and nuclear physicist, Niels Bohr)*

Probabilities and prediction are not the same thing. A great example is rolling a pair of six-sided dice. Assuming the dice and roll aren't fixed in some manner, we can establish with a high degree of confidence that there's approximately a 2.7% probability of rolling snake-eyes, or, in other words, about once every thirty-six rolls. Now, if snake-eyes comes up on the first roll, does that mean our probabilities were wrong? No! There's no rational reason to expect the dice to wait until the thirty-sixth roll to turn up snake-eyes. The key thing to keep in mind is that establishing probabilities is not the same thing as foretelling the future.

The first point many people make about my dice example is that the dice and/or roll may, in fact, be fixed, which can skew the results and invalidate the estimated probabilities. Absolutely true. A few things to consider are:

- All decisions are based on incomplete information. The world is simply too complex to know every conceivable factor, however a rigorous and logical analysis methodology will tend to reduce the number of unknowns.
- Truly significant variables are often either obvious (such as -- “Hey, that die only has five sides!”), or can be fleshed out through thoughtful evaluation.
- Consider the incredible variety of assumptions that go into defining and applying best practices...

By understanding the probabilities surrounding a scenario (even approximately), we’re able to make better informed decisions than if we didn’t consider them at all.

Things change

The argument that a risk assessment done today may be invalidated by changes in the risk landscape tomorrow is absolutely true. Of course, those same changes in the risk landscape are likely to affect how well a best practice is going to perform. The bottom line is that yes, our problem space is dynamic and, yes, regardless of how you choose to approach security you’d better pay attention to how the landscape is changing. It comes with the territory.

Open-ended problems

If we want to understand our problem space better, we’d do well to understand that risk issues are “open-ended” in nature rather than “well-structured”. Well-structured problems can be reasoned to a single correct answer – e.g., $3+3=6$, or “Will I overdraw my bank account if I write this check?” Open-ended problems, on the other hand, are those that can’t be reasoned to a single, undisputed correct answer. Examples of open-ended problems include:

- What’s the right solution for peace in the Middle East?
- What’s the best financial investment or insurance plan?
- Should I press the accelerator or the brake at this yellow traffic signal?

Most of the information security problems we face are open-ended – in other words, there are few clear, undisputed correct answers. Examples of open-ended questions we’re forced to deal with in security include:

- Is this policy exception request acceptable?
- How do we prioritize our many security improvement efforts?
- How much risk does this combination of control weaknesses represent?

Because these issues defy simple, indisputable answers, and because each of our circumstances will vary, we’re forced to apply our judgment and critical thinking skills. Unfortunately, in the absence of a clear methodology we tend to rely on informal analyses and our intuition, which introduces several problems:

- There’s greater opportunity for myth and bias to enter into the equation

- Without applying some sort of consistent and rigorous framework or method to this “analysis” there’s a greater probability of overlooking, undervaluing, or overvaluing key factors, which increases the opportunity for bad conclusions.
- Without a consistent framework for how we rate the variables that go into our analyses and the conclusions that come out of those analyses (one person’s “high risk” issue may be someone else’s “medium risk” issue), the results become inconsistent and far less meaningful.
- We can’t consistently and effectively defend our conclusions to the people who hold the purse strings and who write our paychecks.

Even if you accept the argument that a security program should fundamentally be based on best practices, we’re still faced with decisions like those described above every day. And best practice can’t tell us how much risk a policy exception represents because best practice can’t account for the unique circumstances within each scenario. The fact is, whether we realize it or not and whether we like it or not, those of us who actually have to try to manage risk cost-effectively within an organization are forced to prioritize, make decisions, and defend/explain our rationale within a complex environment. Best practices, by themselves, can’t help us with this.

Regarding “Proof”

The difficulty of establishing “proof” that a security approach has been effective is every bit as true for best practices as for risk management. Both require data, and even the May article stated that there are no statistically relevant data. Consequently, in both cases we’re left with assumptions regarding whether an approach has been effective.

Keep in mind that regardless of how you go about managing security, you’re going to have incidents -- maybe even significant incidents. Recognizing that no security program is going to be perfect at eliminating loss, it’s naive to characterize a significant loss event as clear evidence of an outright failure of the program, regardless of whether that program is risk-based or standards-based. Likewise, in the absence of a significant incident, it’s naive to claim the security program a success based on the management method. Borrowing from the May article, “You can’t prove that something didn’t happen” as a result of security measures that are or aren’t in place, whether they be best practices or something else.

Garbage in, garbage out

A very legitimate concern in analyses of any type is “garbage in, garbage out”. As relates to information security risk, the concern is that without empirical data an analyst may:

- Game the system - i.e., consciously or otherwise choose inputs that result in a desired or expected outcome
- Estimate poorly, due to a lack of experience, understanding, or other cause

Of course, gaming the system is also possible using data, but we can reduce the potential for gaming (with or without data) through peer reviews and by requiring analysts to document the basis for their estimates. Further, it’s even easier to game a system (consciously or otherwise) in the absence of documented analysis.

In my discussions/debates on this issue to-date, the second issue (poor estimates) seems to generate greater heat from those who stand against risk analysis. The phrase I most often hear is

“wild-ass guesses”. Clearly, any estimate is sensitive to the experience and knowledge of the person making the estimate, but a few things to consider are:

- Poor measurements are also possible using data, for example if the data selection criteria or sampling are flawed
- Wildly inaccurate estimates are usually going to generate results that don't pass the laugh test, and a logically structured analysis method allows us to identify and troubleshoot such errors
- The vast majority of professionals I've worked with are capable of generating reasonable estimates, particularly when using a logically structured approach
- Analysis quality can be further improved by having more than one person involved in the analysis whenever possible

A Useful Combination

The fact is, both risk-based management and best practices are important tools for our profession; each having its own strengths and weaknesses. By recognizing and leveraging the strengths of each, we stand a better chance of being successful. Bottom line -- choose the approach that best suits the problem at hand.

Providing Value as a Profession

Any grade school graduate can cite a published standard or compare a checklist against what they see in front of them. As professionals, we provide maximum value to our employers by being able to apply our subject matter expertise as well as critical thinking skills to work toward optimum solutions of complex problems that can't be boiled down to a simple set of rules. Effective information risk management can be complex and difficult – even if you leverage a solid risk analysis framework. This difficulty shouldn't, however, compel us to throw our hands up in defeat and rely solely on attempts at black and white solutions to complex, open-ended problems.

Another point to consider is that we're not only responsible for helping our employers manage risk, we're responsible for helping them manage it as cost-effectively as possible. Unfortunately, strict adherence to a one-size-fits-all best practice approach doesn't enable us to tailor our solutions to our specific circumstances. Sometimes best practice will be the most cost-effective solution for a given circumstance; sometimes it won't. As a responsible professional, the important thing is being able to recognize the difference.