Information Security Risk Assessment- pitfalls and practicalities

Michael Barwise

In an ideal world, information security management would be a simple matter of applying all relevant security measures immediately. But in reality our resources are limited, so we must select a subset of the available options that maximises real security within our budget. We have to prioritise, and the tool is risk assessment. Sadly, risk assessment is the most misunderstood and ill-performed component of information security today, and as a result, corporate security implementations are often extremely poorly focused. I believe there are three primary contributors to this: failure to understand what risk really is, unreliable evaluation criteria, and concentration on technologies at the expense of business issues.

So what is risk? We all use the word in everyday life. The weather forecast mentions a "risk of rain on Tuesday", meaning a possibility. Climbing a long ladder is "risky", meaning dangerous. Neither of these is really risk: strictly, it is a combination of the two. For infosec practitioners, the risk of a breach would be the loss (in, for example, dollars) incurred as the result of the breach multiplied by the probability of that breach causing the loss.

However, the infosec community has invented numerous alternative, more complicated and less precise models of "risk". The most prevalent is to subdivide risk into "threat", "vulnerability" and "likelihood", evaluate each using subjective criteria, and then combine them in some empirical way to arrive at a notional value for risk. Threat and vulnerability are typically evaluated intuitively using verbal hazard scales such as "low, medium, high". Because of their subjectivity, these categories are extremely difficult to assign to threats or vulnerabilities, or indeed, to interpret, with any degree of confidence.

Such uncertain classifications are a major contributor to the failure of infosec risk assessments, but we continue to use them because they are convenient, and we are unaware how subjective our decision making is. Although the quality of even expert judgement has been widely researched in the context of large-scale capital projects, and found to be on the whole quite poor, infosec professionals have not yet taken these findings on board. Predictably, pretty well all the published infosec risk management guidance stops short of describing how you actually determine the values you need to use in these "risk models".

Likelihood is generally evaluated using statements such as "twice a week" or "once in three years", which lead us to confuse statistical probability with the realities of event occurrence in the operational context. A general lack of statistical expertise among practitioners causes us to miss the fact that statistical probability is not the whole story in the context of infosec.

Basically, there are two possible scenarios: either a given breach occurs or it doesn't. This is a bit like tossing a coin. Given a fair (perfectly balanced) coin, over a large number of tosses we would expect heads to come up half the time (a probability of 0.5). But this probability says absolutely nothing about whether heads will come up next. Equivalently, a probability of a given infosec breach occurring "once in five years" does not mean it won't happen twice next Tuesday. And there is another very real problem. Unlike our coin, which does not change its fairness while we gamble, infosec threats and vulnerabilities are constantly changing in both nature and prevalence. If we want to use the likelihood of, say, a threat as a component of our prioritisation strategy we will need, in addition to knowing its overall statistical probability (supposing this can be determined), a large amount of reliable evidence of past occurrences to determine how its successes and failures have been distributed across targets and in time. Unfortunately, we just don't have this for most threats and vulnerabilities, mainly because the community as a whole has not kept good enough records. Even in the case of Code Red, which is relatively well documented, there is less and less information as time passes and attack patterns decrease in intensity.

Even if this were not the case, using technical threats and vulnerabilities as the sole basis for risk assessment is not satisfactory. On the one hand, one can get bogged down in comparisons of attacks and their fixes, and on the other, it is easy to fall into the trap of thinking in crude terms such as "firewall security" or "anti-virus", forgetting the business purposes for which the technologies are deployed. This often leads to extreme precautions in a given technical sphere which are completely undermined from the business perspective by omissions in another. Ultimately, it is the business

information assets we are trying to secure: not the “perimeter” or the “server”, so we need to know where business value is concentrated.

So, we need an alternative approach to risk assessment that does not suffer from these limitations. I believe we must abandon technocentric risk models based on threat-likelihood, with all their uncertainties, and instead establish business-oriented security priorities on the basis of some relatively solid criterion. Only once this has been done should we start to investigate the technical issues relevant to our identified priorities.

The first step is to identify our information assets by examining all our business processes, and to determine how each information asset is handled at every stage of each business process. The media and infrastructure components involved in those processes must be identified and documented (including the phone calls, faxes and post-it notes: they’re relevant to information security too).

Now we can assign a value to each information asset by establishing the financial loss that would result from a total breach of the asset in the context of the given business process. We can categorise into quite broad bands (I generally advise the use of five bands), but should always call on our legal and insurance advisors to check our evaluations. Having done this, we will be able to identify which information assets represent our greatest potential losses. These are our highest priorities.

The next stage is to map the prioritised information assets back to the media and infrastructure components already identified, so we can calculate the aggregate value of information relevant to each. Only than do we start to think “technical” and investigate the threats that target them. Remember though, we’re not just talking about “hackers and viruses”: more information is jeopardised daily by bad business procedures than by all the Internet threats combined.

Once the mapping is complete, we can prioritise security measures on the basis of aggregate potential losses for groups of assets that map to securable entities such as LAN segments, media and business units. So you see I am not suggesting you ignore technical issues: just that they are most appropriately dealt with in the second round, not the first. This method will work. The only question is whether we can strictly speaking call it “risk” assessment. I guess it’s really “requirements analysis”.

Finally, this is a programme, not a project. There must be a continuous process of incremental review and improvement. You have to keep your prioritisation database up to date, otherwise it will cease to be reliable. And you won’t get it perfect first time. But because the process is methodical and based on facts, you can adjust it globally to correct errors as you identify them. Alternatively, you could always save yourself the effort by buying some off-the-shelf risk management software and ticking the boxes. It’ll output some pretty reports. But how will you know it’s delivering real security where it’s needed?

Originally appeared in Computer Weekly, October 2002