Risk judgement and knowledge

Michael Barwise

When discussing information risk, most of us seem to concentrate on the sharp technical end where things change rapidly. But if we step back a bit and consider even technology risks from a business perspective, they become more generic and stable. So for example the risk of data leakage due to a particular piece of hot malware is very difficult to evaluate reliably and can change literally overnight, but the risk of data leakage due to the corporation's general posture in respect of data and systems management is more reliably determinable and consistent over time. Of course it's necessary to consider your risks in both business and technical terms, but the conventional approach to treatment is very different at the two levels. If we allow the technical position to dominate our thinking, we are forced into a reactive mode: waiting for signs of intrusion and responding to specifics. If, however, we start from a business perspective and work downwards towards the technical, we can adopt a proactive stance resulting in resilience against generic threats. Although we can never entirely eliminate the unexpected, I believe the latter is the better starting point as it can reduce its likelihood significantly.

Many businesses are much more interested in certification to a recognised security standard than they are in the details of the security it's supposed to represent. It's obvious why - you can't process credit cards unless you've passed a PCI-DSS audit, and you can't land government contracts in many jurisdictions without ISO 27001 certification. So these are business enablers. It's therefore not surprising or even unreasonable that we tend to look on these standards as "gold".

But it's important not to confuse the two purposes - getting certified does not ensure you're secure. There's no established standard used in infosec that goes into enough detail to ensure genuine best practice (rather than just most common practice) in the decision processes that underpin security. The standards we have only prescribe processes that must be performed in organising risk decision making. So compliance may or may not indicate you've made good decisions - it just shows you've made them in a standard way, not that they were optimum (or even adequate).

Looking at other spheres where standards apply, we find rather a different picture. For example BS 5701-4:2003 *"Guide to quality control and performance improvement using qualitative (attribute) data - Part 4: Attribute inspection performance control and improvement"*. To quote the abstract *"...This standard deals with measuring and improving the quality of decision-making in the classification process itself when subjective judgments are involved such as determining whether a particular flaw or imperfection is present or not..."*

We could well make use of this kind of detailed standard in our infosec risk judgement processes. However, none of the accepted information security standards go into this level of detail. Indeed I have noted a gradual decline in the detail provided by standards in general over the years. For example, the now withdrawn BS 600:1935 "The application of statistical methods to industrial standardization and quality control" was even more detailed - sufficiently so to be a useful pedagogical document for the teaching of applied statistics.

So to reiterate, the whole issue of risk management comes down to the quality and consistency of your judgements, and that's not covered by any standards we currently use in infosec. We have to go elsewhere for this knowledge, but without it compliance with standards is mere theatre. There's a bunch of psychological traps we fall into if unaware (so we must become aware) and there's a body of math that must be understood, before any of us have a hope of making consistent reliable risk judgements. If we don't acquire this basic knowledge base before we embark on risk judgement we're just crystal gazing.

Circulated as a white paper, December 2012