Information Security Now – 39

John Mitchell

Regular readers of this column will know that the risk equation has two elements: likelihood (probability) and consequence (impact). They will also be aware that these two elements may be mitigated by the application of control(s). The likelihood element is usually measured using some form of probability scale, whilst the consequence tends to be measured by financial impact. The two elements are then multiplied together to predict the hit on the enterprise's bottom line. I accept that this is being expressed in simplistic terms, but when you only have a thousand, or so words to play with, it's the best I can do. The mitigation of each element by the implementation of some form of control implies that there is a raw state where no controls are in place (inherent risk) and a managed state which assumes that the controls are effectively operating (residual risk). Residual risk does not imply no risk, but ultimately it is what the enterprise agrees to live with and ultimately becomes the retained risk for the enterprise. This retained risk may be known to the Board whom accept it, or it may remain unknown by them and accepted by default. This latter situation is very common for IT risks, were the Board often abdicates its responsibility on the basis that it is either too difficult to understand, or cannot fully comprehend the impact. This is of course the responsibility of the CIO who must clearly explain the consequences of an attack in business terms. Statements such as 'Port 80 on the firewall is insecure', may be technically correct, but cut little ice with the Board.

For this retained risk to be meaningful it is obvious that both elements of the risk equation are accurate and here comes the problem. The likelihood is basically an informed guess and consequence must include the full impact if the risk crystallises. When we consider cyber-attacks it would be safer to assume the likelihood is absolute at the inherent risk level and still pretty high at residual risk and before you start shouting 'firewall' and 'anti-malware' at me please remember all of the companies that have suffered embarrassing cyber-attacks in the last few years. They had all of those things, but the attack still succeeded: centrifuges were persuaded to run at a damaging speed; state secrets were released to the media and customer financial data was stolen. So it may be prudent to assume that an attack will be successful. The question then becomes what is the impact and how do we handle it? The impact of bad publicity is difficult to estimate, but we must at least consider the potential loss of customers, a drop in the share price for a listed company and potential regulatory sanctions. The problem then becomes one of damage limitation. Talk-talk, the telecoms company received more stick for not quickly owning up to the theft of customer data, than it did for the scale of the problem (which ultimately turned out to be less than originally estimated). So careful handling of the media is at least as important, perhaps more so, than the technical response.

I conduct knowledge management audits for my clients. Without fail, these show that in the event of a major incident (not just IT) the most important person in the enterprise is not the CEO, nor any other person in the 'C suite', but rather the media relations person, because they are the point of contact with the media; either providing information via press releases, or responding to queries from journalists. IT journalists tend to be knowledgeable of the technology and will quickly identify errors, or omissions, in information and will react accordingly. Cover-ups are usually quickly identified and harshly treated. Waffling is treated with disdain. So it is sensible to have a prepared response to the most likely scenarios where you just need to fill in the blanks when the bad thing happens.

The standard scenarios are a DDoS attack (no loss of data); loss of company data (but no loss of customer data); loss of customer data (non-financial); loss of customer financial data. These scenarios can be flexed to cover external attack, internal attack and whether initial disclosure came from you, or someone else. The US government was extremely embarrassed when Edward Snowden released his first tranche of stolen data because they didn't know until then that it had been stolen. This is an important point in any response plan, because data theft usually involves copying of the data you still have the original, so may be unaware that you have been attacked until you read about it. This of course is very embarrassing, as you are on the back-foot from the start and are trying to handle the media whilst trying to find out what has happened and when. The when is really important because it may provide an indication of the scale of the problem. All the more reason to have a prepared response.

At the beginning of this article I mentioned the move from inherent to residual risk. This is achieved by the implementation of controls and provides a control line which is capable of measurement. So I spend a lot of my time evaluating risk registers with particular attention to this control line. If the controls are well designed, effectively implemented and regularly monitored then the residual risk estimate is likely to be correct. However, this is seldom the case and I despair at the easy acceptance of the predicted residual risk by senior management. They will argue that they are not IT experts and have to rely on their IT colleagues. I have little sympathy with this view. They should ask for an independent and objective appraisal by a qualified IT risk professional. Also, as I have explained, much of the response to an attack is handled outside of the IT function and that person is not an IT expert either. Indeed, this provides a perfect example of a risk event arising in one function with the consequence being handle by another. Which is why any company needs a joined-up risk management process which goes horizontally across functions. This may require one function responsible for managing event probability and another for dealing with the consequence fall-out.

John is Managing Director of LHS Business Control, a corporate governance consultancy. He is a member of BCS Council, it's Audit & Risk Committee and Chair of the Information Risk Management and Assurance (IRMA) specialist group. He can be contacted at: john@lhscontrol.com, www.lhscontrol.com, or on +44 (0)7774 145638