Information Security Now – 33

I spend a lot of time reviewing risk registers. It is an amusing adjunct to my job as an IS auditor. 'Amusing', I hear you say. 'How can something so serious be amusing'? Well, it 's the law of unintended consequences. The three things on a risk register which often cause me to chuckle are: the inherent risk score; the controls; the residual risk score. Why the amusement? Primarily, because of the optimism of the creators of these important pieces of information. Let me explain each in turn. The inherent (raw, or gross risk) is where you would be without any controls in place. It comprises two components: likelihood (possibility) and consequence (impact). So if you were (say) a large on-line auction house assessing the likelihood and consequence of an unauthorised person stealing your customer database, then without any controls in place you would likely score the equation as high likelihood and high consequence. If you used a red/amber/green (RAG) status it would be red/red. You would probably assess this as undesirable and decide to put some control(s) in place. Now the risk equation is remarkably fickle and often you find you can only manage one side of it. In this case you could probably reduce the likelihood side of the equation by using some form of access control and privilege allocation. Indeed, you may decide this is so good that you reduce the likelihood of unauthorised access to low (green). But what about the consequence if unauthorised access is obtained? Well, it is still disastrous and should be scored as high (red). So the score has changed from red/red to green/red. Which is still pretty frightening, but as you have dealt with one side of the mess you convince your superiors (if they are even interested) that you have reduced the risk. Even more so if you make the mistake of multiplying the two attributes together, which many risk charlatans do. Here is an example. Let's assume that we have a range of one to five for each attribute. In the original no control (inherent) scenario, we score each attribute as five and multiply them together to give an inherent risk score of twenty-five. After putting in our access control we now rescore the likelihood as one, but the consequence remains at five. Multiply one by the other and our risk score is now five, an apparent five-fold reduction in risk. What a result! However, a low likelihood is not a 'no' likelihood and if our access control is breached we are in serious trouble. However, using the multiplication mechanism it does not look that bad. After all, it's only a five.

The introduction of the access control has reduced the likelihood of a breach from red to green, but then only if the control is one-hundred percent effective. This is where the skill of control evaluation comes in and is this component which causes me so much amusement. In the case of the eBay breach we know that an internal employees' access credentials were breached. Once 'they' have your access credentials, then they have your privileges. They effectively become you and no amount of intruder detection is going to prevent them from doing everything that you are allowed to do. No alarms are triggered; just you doing your job. Which is why it took a couple of months for the breach to be noticed. Now it is a dichotomy to me that organisations appear to have different authentication criteria for internal and external access. For the former it is usually a simple user ID and password, while for the latter it is often a one-time password generator. I know a number of banks where this holds true and have never figured out why they discriminate between the two;

especially when internal staff often have greater privileges than external users. Breaches occur because of a combination of complacency and trust. Neither of which are a control. If we assume that the eBay breach was not conducted by an insider (and we are told that this was the case), then the attacker gained the access credentials of a privileged staff member. If a couple of simple authentication factors, say the one-time password generator with a token, had been a requirement, then the attack would have been thwarted at birth. Truly moving the likelihood from red to green. You notice that it still does nothing to lower the consequence which remain red.

I use a simple pseudo-mathematical mechanism to score control effectiveness for both likelihood and consequence, which I will not elaborate on here due to the word count imposed by the editor. I use this on every so-called control in the risk register to see if the risk is really mitigated by the control. The answer is usually depressing to the risk owner who often asks 'what else can I do'? The answer is to employ a control expert (beware of charlatans). S(he) may depress you even more, but at least you will truly know the risks that you a living with. On a more positive note the resulting dialogue often raises both risk awareness and control effectiveness. IT people tend to be the optimistic Tiggers from Winnie the Pooh, whereas us IS auditors are the pessimistic Eyhaws. However, unlike Eyhaw we have some pretty good tools to support our views on the effectiveness of your controls.

John Mitchell

John is Managing Director of LHS Business Control, a corporate governance consultancy. He is a member of Council 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