Information Security Now – 45

Denial of Service
(From WannaCry to ReallyCry[1])

John Mitchell

The Institute dedicated much of its Winter 2016 edition of IT Now to Ransomware, which was well before the recent NHS attacks, so it was unfortunate that the powers in the Boardroom did not have it drawn to their attention before things went belly-up. Had they been so informed, then the outcome may have been a little less severe. You may well then be asking why have I returned to a topic which has received so much publicity? Locking the door after the horse has bolted – again? Well not quite. Ransomware is simply another denial of service problem, together with power outages, DDos attacks, general malware, poor software integrity and data corruption. Denial of service can be divided into the further categories of deliberate and accidental. By their very nature, ransomware, DDos and general malware are deliberate, whereas power-outages, software and data corruption may be either deliberate, or accidental.

So why was WannaCry so successful in its propagation? After all, most entities have anti-malware protection. According to Malwarebytes Lab, Wannacry hunted down vulnerable public facing service message block (SMB) ports and then used the NSA-leaked EternalBlue exploit to enter the network and then used the DoublePulsar exploit to establish persistence and allow for the installation of the WannaCry Ransomware. So, it was not spread by email and any protection relied on the installation of a patch to prevent entry. Microsoft even released a rare emergency patch to help protect Windows XP devices (the company hasn't officially supported XP since 2014), but if the patch was not installed, then WannaCry had free reign. The story of its killing is well known and I will not repeat it here save to say that it appears that the kill switch was intentionally built into the ransomware to stop its propagation if so required.

Using basic risk analysis, we can establish that the consequence from all the denial of service events mentioned earlier is indeed service unavailability. However, the cause (event) for each is markedly different. In ISO 27000 terms we have multiple threats leading to the same consequence, but the treatment for dealing with these threats is likely be different for each one. Those readers familiar with risk analysis will recognise the standard risk equation of likelihood and consequence, or probability and impact, depending on your preferred nomenclature. The standard way of calculating how much you should be spending on managing a particular threat is to multiply the likelihood (from zero through one) by the financial impact. Thus, if the likelihood is 0.5 and the estimated financial impact is £100,000, then you should never spend more than £50,000 on managing the threat. The challenge here is two-fold. First, our predicated likelihood is simply a guess (we do not have the guidance from the

[1] The fall-out, after the event has crystallised, from the 'I told you so' predictive syndrome.

morbidity tables of the life insurance companies) and second, the value placed on the consequence must include all costs, both direct and indirect, incurred by the business if the risk crystallises. Both of these challenges are non-trivial, as the not-for-profit NHS and the for-profit BA found out to their respective costs.

The reports that the intensity of the Grenville tower block inferno may be partially as a result of saving £300,000 on a total refurbishment cost of £10 million is an example of reducing security spend without estimating the full cost of the impact in the event of risk crystallisation. I suspect that £300,000 is going to appear to be small beer in the total remediation costs and will not even pay for the forthcoming public enquiry.

The challenge in calculating how much you should spend in to achieve your availability objectives is obscured by the way we view the IT budget. IT is universally viewed as a cost to the business, with the constant wrangling between the CIO and CFO over what it should be. However, if we examine IT from a value, rather than a cost perspective, then the value of IT to the business is often the total value of the business. This is because many modern businesses would be unable to function without IT. If a denial of service problem stops the business in its tracks and there is no possibility of recovery, then the business dies. If we have recovery, then the loss to the business is the full-cost from the time of service interruption to service recovery. In the case of BA an estimated cool £100 million. I have some fancy pictures to illustrate this, but as the editor strictly enforces space allocation, contact me directly if you wish to see them.

Businesses seldom die as a result of a service interruption, but they are often woefully incompetent when it comes to estimating the full-cost of the interruption; especially when considering the impact on their brand value. This is important because it has a direct impact on how much you should spend on security. Unfortunately, IT security, which includes service availability, is usually viewed solely as an IT matter and all costs are estimated within the IT arena. Thus, security spend is seen as a component of the IT budget, rather than the total business operating budget. A recent Gartner survey[2] indicates IT security spend as being between 1% - 12% of the *IT budget*, depending on industry sector and entity size. You will notice that the survey deals with security spend as a percentage of *IT budget*. However, (and this is my Eeyore moment) because IT impacts on the entire business I consider it unfair for IT to bear the entire IT security burden from within its limited budget. Especially so with regard to defending against attack software, where, with the best will in the world, the malware writers are always going to be ahead of the defenders because it is virtually impossible to defend against something you don't yet know about. As Karl Popper once wrote, 'You don't know what you don't know until you know it', which is likely to be too late. So, if we can't defend against it, then we should assume that the probability of denial of service in the future is pretty much certain. Therefore, we need to spend money on fast detection and recovery in order to minimise the service interruption and therefore the cost to the business.

[2] IT Security Spending Trends SANS Institute - 2016

Every organisation needs an Eeyore, the pessimistic donkey in A. A. Milnes Winnie the Poo series of books. As an assurance provider, I am frequently accused of being an Eeyore and not being a team player. I take that as a compliment, because the danger of the 'consensus' ideology is that alternate views are not tolerated. So, this is my ReallyCry moment. I cannot remember the number of times when I have been told the following after I have pointed out the dangers of denial of service, whether accidental, or deliberate. First, it hasn't happened yet and therefore by implication it never will. Second, the mitigation cost is too high; implying that they really know the full business cost of an interruption, when they probably don't. Third, our current security strategy is adequate; implying that they have effective business continuity/disaster recovery for all future threats, even though they don't know what these are going to be. These three answers really do want to make me cry.

John is Managing Director of LHS Business Control, a corporate governance consultancy. He is a previous member of Council and the Risk, Audit and Finance Committee. He is currently Workshop Liaison for 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