Risk Judgement Quality and the Security of Critical Infrastructure

Michael Barwise

The first challenge faced when attempting to secure critical infrastructure is the accurate definition of "critical infrastructure" itself. In a densely interconnected information sharing community where partner access to networks that would once have been isolated has become the norm, I feel this is not trivial to accomplish. It needs a combination of broad technical knowledge and a sound business-oriented appreciation of the information flows that the infrastructure supports. Without this broad perspective, we are likely either to concentrate on securing only the obvious major physical components of a common infrastructure at the expense of less visible but equally vulnerable potential points of breach (particularly at partner interfaces), or to run the risk of inhibiting the fulfilment of business objectives by over-zealously locking down technologies. As Ruane[1] has suggested, *"identify the wrong mission-critical asset and [there is] a good chance of wasting time, resources and manpower ... while a true (but unknown) mission critical asset remains vulnerable ..."*

The second challenge is the proper identification of vulnerabilities and their real proximate causes. Taking software vulnerabilities as an example, the scale of the problem is not primarily due to the cleverness of the "hacker" community, but to the intrinsic insecurity of current software itself. This is unlikely to be rectified in the near future unless there is a radical shift in the perceptions of developers,[2,3] and it forces the deployer and user into an inherently weak reactive position, with the obligation to apply fixes as they are released in response to recognised specific threats. We have to recognise that our fundamental vulnerability here is the use of reactive response as a primary mode of defence, and not any specific "cyber threat". I concur wholeheartedly with Richard A. Clarke, Special Advisor to the US President for Cyber Security[4] "*It doesn't matter whether it's al Qaeda or a nation-state or the teenage kid up the street, who does the damage to you is far less important than the fact that damage can be done. You've got to focus on your vulnerability and not wait for the FBI to tell you that al Qaeda has you in its sights.*" A better defensive strategy would in principle be proactive vulnerability assessment by infrastructure owners and managers at a higher conceptual level rather than an *ad hoc* technical level. This is, however, only one side of the coin. Alone, it will not prevent security management remaining a permanent and significant drain on resources. We also need to address the root causes of insecurity at source. Although Ross Anderson[5] and others have clearly expressed the intractable nature and scale of the problems facing software developers, I believe that there is nevertheless a potential long-term benefit in persuading them to tighten up on the intrinsic security of their products. It may indeed require a new way of thinking on their part. Research I have conducted suggests that there are fewer than half a dozen technical proximate causes underlying all software security vulnerabilities. Should this be accepted, the problem would become manageable by concentration on this small set of real issues, rather than on the dishearteningly vast array of their individual implementational manifestations. A lead must be taken in this respect for any real long-term improvement to be made in critical infrastructure security, and, in the words of David Blunkett[6] in order to lead *"you must make a decision and take people with you. You can't just go for the lowest common denominator."* There is clearly room here for Government to take the initiative when commissioning IT, although I have so far found most agencies loath to allocate the requisite time and resources to proactive security assessment.

Malicious software breaches and "hacking" are hot topics (particularly in the face of the "cyber terrorism" spectre). They also genuinely represent a clear and present threat, particularly in the context of control infrastructure (e.g. industrial control and signalling). They are nevertheless only part of a much larger picture. I believe we must take a broader view of critical infrastructure security that considers information assurance as a whole, rather than concentrating solely on covering emerging technical vulnerabilities in a reactive and *ad hoc* fashion. Over the last few years, there has been a steady trickle of unintentional information leakage from both public and private sector information

[1] Ruane, LtCol T. Operationalizing Critical Infrastructure Protection. IATAC IA Newsletter, 5:3, 2002. p17.

[2] Anderson, R. Why Information Security is Hard. UCCL. 2001.

[3] Brady, RM *et al*. Murphy's Law, the Fitness of Evolving Species, and the Limits of Software reliability. UCCL Technical Report. 1999.

[4] quoted in: Ruane, *op cit*.

[5] Anderson, R. Why Information Security is Hard - An Economic Perspective. University of Cambridge 2001.

[6] Blunkett, D. Interview, "Today". Radio 4. 04/02/2003.

systems[7] independently of any hacking activity, and such leakage renders the business supported by critical infrastructure just as vulnerable as many directed attacks on the infrastructure itself. As Gene Spafford recently testified before the US Senate on behalf of CRA/USACM[8] *"... those of us working in the field have learned that the issues are really larger than simply computer security. Information assurance covers issues of building safe and reliable information systems that are able to weather untoward events no matter what the cause- whether natural disaster or caused by a malicious individual."*

I believe a prerequisite for proportionate risk management is the capacity to evaluate potential threats within this broader framework, rather than solely in the context of their *prima facie* technical attributes or of arbitrary topical concerns. To do this requires, in addition to specific expertise in a range of technical, business and numerate disciplines, a specific state of mind which is not innate. As Morgan and Henrion[9] among others have shown, expert judgement needs to be elicited under carefully controlled conditions to maximise its objectivity. This is all the more the case where decision-makers are less than maximally expert: a common circumstance in the IT employment space, due to the limited availability, and cost of retention, of leading experts. We ultimately need infosec practitioners in the front line to think in a way that is rigorous in respect of detail and simultaneously broadly perceptive, and which allows them to establish fundamental premises that can subsequently be used to build on experience. This was encapsulated superbly by John Dewey[10] as long ago as the early 1920s *"... men do not begin thinking with premises. They begin with some complicated and confused case ... Premises only gradually emerge from analysis of the total situation. The problem is not to draw a conclusion from the premises ... [but] to find statements of general principle and of particular fact that are worthy to serve as premises."*

My experience of information security risk evaluation in both public and private sectors has shown me that, in most cases, the requisite analysis of the total situation does not take place, so evaluation rarely proceeds beyond Dewey's complicated and confused case. Evaluation processes are generally diffuse and *ad hoc*, are based on assertion without proper evidential support, and operate without consistent controls. Often the documentation of risk evaluation is rigorously controlled, while the evaluation process itself is not, resulting in a paper trail that gives a superficial but unjustified impression of robustness. To combat this situation, I believe that, in addition to the short-term need to inform the owners and operators of critical infrastructure of specific threats and countermeasures, there is a long-term obligation to develop robust consistent evaluation methods that can be used by all parties to the support of critical infrastructure. In developing suitable risk evaluation methods, we must explicitly seek to close the current divide between the high-level view of the business risk management community and the technocentric perspective of the front line IT security practitioners upon whom we ultimately depend for the physical implementation of critical infrastructure security. There are, for example, numerous high-level standards and guidelines that address business and information technology risk assessment, but although they identify objectives and desired results, none of them delineate proven methods or procedures whereby these can be reliably accomplished. Thus they cannot cater for the very common situation where inexperienced staff may be charged with ultimate responsibility for entire mission critical security risk assessments by virtue of their job title rather than their expertise. Until the general level of expertise improves, this large and critically placed echelon must be given robust processes to follow, documented if necessary to a procedural level. Most IT practitioners in the front line cannot be expected (for example) to decide spontaneously on the probability of a rare event with any reasonable degree of confidence. They need to be told what basis, and what kind of enquiry, to use, and most importantly, how to validate the judgements they make.

The third, and probably the greatest, challenge we face, therefore, in securing critical infrastructure is not the solving of technical problems, nor is it even the timely recognition of present threats. It is a lack of expertise in robust risk evaluation methods where it is most needed: in the front line. Unless this can be addressed, the only prospect for agencies such as the US-CERT or the UK NISCC is to deliver a perpetual stream of technical alerts, and for those tasked with business information security to continuously and reactively respond to them. This creates an unacceptably fragile dependency on a

[7] e.g. UK Treasury. Budget 2002. MS Word drafts including change history, published on the web.

[8] Spafford, E. 2002. http://www.cra.org/govaffairs/advocacy/spaf_testimony.html

[9] Morgan, MG *and* Henrion M. Uncertainty. Cambridge University Press. 1998 (6th imp.).

[10] quoted in: Gaskins, RH. Burdens of Proof in Modern Discourse. Yale University Press. 1992. p144.

single point of failure, and it will indeed fail, as has been shown by the increasing number of "zero-day vulnerabilities", that is, threats that emerge too swiftly to be officially reported. In the idiom of the Chinese proverb, we can either give a man a fish every day for his whole lifetime, or we can supply him with food while we teach him to fish. Clearly, the latter is the better option. However, the fundamental prerequisite is recognition that a problem exists to be solved. At present, I find this broadly lacking. The quality of information security risk judgement, although in general still demonstrably poor, is taken for granted.

Submission to EURIM, August 2005