

Procrustean Risk Management – Why Critical Systems are at Risk

a report by

Dr Kurt J Snapper

Member of the Security Affairs Support Association, Armed Forces Communications and Electronic Association, and Information Technology Association of America

Dr Kurt J Snapper is a member of the Security Affairs Support Association, National Defense Industrial Association, Armed Forces Communications and Electronic Association and Information Technology Association of America. He is President of ManTech Security Technology Corporation and serves as an independent security advisor for private sector customers. He provides security and risk-management research and development and consulting support to several US agencies and private companies. Dr Snapper's particular area of interest is the protection of critical resources against insiders and sophisticated adversaries in high-risk environments.

Computer and communications security has been a growth industry for more than two decades and the forecast is for strong, continued growth. Both government and industry are expected to continue making significant investments in security products and services. Anecdotal evidence suggests that the cybersecurity profession has not been particularly effective, if success means actually securing the cyber-infrastructure. On balance, cybersecurity vulnerabilities and associated risks have tended to increase, even as the cybersecurity industry has flourished.

The increase in risk exposure is real. Many cybersecurity professionals have noted that the overall cybersecurity exposure of many agencies and companies has not improved much and, in some cases, has declined. Some successes can be identified, but there are also highly sensitive or critical systems that are protected unsatisfactorily. The tally is discouraging. Progress in securing relatively mundane systems is more than offset by the huge risk exposure in critical cybersystems.

As pervasive as problems such as inadequately trained security administration staff and tepid management support may be, these are only indicators of underlying problems. Technologies with known vulnerabilities are being deployed with security architectures that are inadequate in design and cannot be supported or sustained.

This is particularly true for critical cyber-infrastructure systems. The overall degree of exposure has reached the point where the protection of critical infrastructures is vital to industry and government, and the susceptibility of these infrastructures to cyberterrorism and warfare is a growing concern.

Overall risk exposure for critical systems has become unacceptably high because of the inherent vulnerabilities of cybersystems and the growing sophistication, aggressiveness and malevolence of adversaries. That argument is not terribly persuasive because risk mitigation strategies are usually evident but are not implemented. Deliberate decisions to accept high levels of risk are

being made and the problem stems directly from the risk-management process.

The problem is twofold. Senior managers may regard 'risk management' as a rationale for implementing systems with known vulnerabilities because they regard security as a costly nuisance. They often refer dismissively to more robust security requirements as 'risk avoidance', thereby frustrating meaningful consideration of risks and exposure to critical assets.

Procrustean reasoning by security and risk-management professionals themselves is the second and greater problem. In mythology, Procrustes cut and stretched his victims to fit his bed, and, in reality, cybersecurity professionals tend towards 'one-size-fits-all' solutions. Generally, these solutions are adequate for non-critical systems and unsophisticated adversaries, but leave critical, high-value systems vulnerable to attack by sophisticated adversaries.

At the root is the tendency to underestimate the threat against critical, high-value targets. One reason for this is a strong bias that sophisticated adversaries are interested only in classified national security information or an assumption that the probability of an attack is negligible. There also seems to be a pervasive tendency to underestimate the capability and determination of some groups and their ability to circumvent commonplace cybersecurity controls. It is common for systems to be deployed knowingly with vulnerabilities, under the assumption that the probability of an attack is low, with the result that many or most of the critical systems are at high risk.

The perception of risk is driven largely by the number and characteristics of events that have been recorded actuarially or are governed by 'available' documented scenarios. The risk manager is hard pressed to point to sophisticated attacks justifying expenditures for much more than protection against hackers and perhaps the relatively benign insider. Invariably, this leads to an overemphasis on relatively unsophisticated hacker and insider attacks and security protection programmes that are inadequate to sophisticated attacks.

This point is similar to the one that was made by Donn Parker.¹ Parker argues that the profiling of adversaries in terms of characteristics, skills and knowledge can be useful, but is dangerous if people rely on the ‘average’ profile. He argues that risk can hardly be managed because it derives from unknown perpetrators with unknown capabilities and unknown motives.

With regard to critical systems, the main concern is those with the skills, resources and commitment to conduct sophisticated operations to penetrate, attack or exploit networks and interconnected systems. Given the low frequency of attacks against critical systems, the probability of an attack against any specific target is extremely low.

Sophisticated, motivated adversary groups exist and they are capable of defeating the controls in nearly all protection programmes. Since the consequence of a successful attack against a critical system is unacceptable, some recent policies have eschewed threat prediction and, instead, have required protection against a defined threat level. This approach involves establishing a protection level that is commensurate with the ‘attack level’ of the most sophisticated adversary that might attack the system. From a policy perspective, the advantage is that it ensures that critical systems are protected adequately against the high end of the threat spectrum.

Arguably, this approach is conservative, but, given the tendency to underestimate or, more accurately, to ignore the known threat by assuming that “it can’t happen to us”, the conservative approach provides a useful check on other types of risk assessment techniques. It can at least help to avoid the spurious reasoning that is based on underestimates of the threat.

A specific way in which to implement this approach and to avoid the most egregious forms of false reasoning in risk management is to define risk levels and to state that, as a matter of policy, the ‘level’ at which systems will be protected depends on the value and attractiveness of assets. One US agency defines four levels of threat that reflect the sophistication of the adversary generally.² *Table 1* is adapted liberally from the referenced source document. ‘Infrastructure attacks’ is listed in the reference as requiring skills over and above those that are required for attack levels three and four and is therefore shown in *Table 1* as a fifth level.

These threat levels are defined by specific assumptions about the capabilities of adversaries. Protection requirements must still be established that are based on several factors, including the value of assets and their attractiveness as a target for adversaries. It is

Table 1: Characteristics of Adversaries that are Capable of Different Levels of Attack

Attack Level	Name of Attack	Characteristics
1	Unstructured attacks	Basic-level attacks exploiting open doors.
2	Semi-structured attacks	Individuals and organisations using moderately sophisticated methods. Penetration of systems protected weakly, use of automated tools, use of multiple systems to confuse traceback, hiding of activities in penetrated system.
3	Structured information system/network exploit	Substantial organisation and structure of a nation-state, non-nation-state or organisation mounting a concerted effort. Capable of stealing passwords protected by encryption, establishing avenues for future attacks, establishing links from a higher level of protection to lower levels.
4	Structured information system/network attack	Above, plus substantial intelligence resources to plan attack and expertise to conduct the attack.
5	Infrastructure attacks	Level three or four attack capability, plus additional intelligence resources and very sophisticated skills to mount an attack against an infrastructure system.

useful to consider the realistic level of effort and sophistication that an adversary would use to attack a given system and to design a protective strategy to counter that level of threat.

It is also useful to postulate an attack level – as defined in *Table 1* – and then ‘design’ the most cost-effective strategy with which to counter that threat. For example, if an attack level four is postulated, the system would need to be designed to prevent or defeat an attack from a sophisticated and determined adversary.

Clearly, there is uncertainty about adversaries and the attack level that they represent, but some reasonable assumptions can be made about them and the targets that they would find sufficiently attractive to dedicate considerable time and resources to in order to conduct an attack. For example, it is plausible to assume that sophisticated adversaries would be attracted to high-value targets and it is therefore reasonable to attempt to develop cost-effective protection.

Once assumptions about the attack level are made, the development of cost-effective strategies can be a demanding exercise for attack levels three to five. In the case of adversary groups that are hoping to conduct an infrastructure attack or to engage in the long-term exploitation of a sensitive system, the capabilities that the adversary group has at its disposal must be considered.

1. D Parker, “Understanding ‘Peopleware’”, Information Security, June 2001.
 2. US Department of Energy, Cyber Security Threats, Version 1.2.

Table 2: Design Basis Threat for Primary System Components

Type of System	Design Basis Threat	Comments
Worldwide corporate network	2	<p>Important for ensuring availability of communications.</p> <p>Contains no sensitive information, but is relied on for transmitting sensitive information.</p> <p>Reasonable precautions against denial of service required.</p>
Sensitive outer networks	2	<p>Used for the conduct of routine agency business.</p> <p>May contain personnel and other sensitive information that requires a reasonable degree of protection.</p>
Critical sensitive inner networks	4	<p>Contains information that is critical to the conduct of organisational mission.</p> <p>Requires cyber, physical and technical protection against a determined adversary.</p>

This methodology is being used for the development of security architecture for the US foreign affairs community. Several committees and panels have indicated the need for modernisation of the community's computing and communications infrastructure. Modernisation is one of the major congressional concerns. The protection of these systems is also of concern to congress and senior managers within the foreign affairs community.

Senior risk managers and security staff recognise the importance of prudent measures to protect sensitive, unclassified information and to ensure the continued availability of critical systems and communications. For example, Internet Web browsing could expose these systems to attack from a wide range of relatively sophisticated adversaries, resulting in a compromise of sensitive information or denial of service.

The modernisation process, therefore, was initiated along with the development and implementation of a risk mitigation strategy. The objectives were to correct any deficiencies in the existing systems – mitigating any vulnerabilities that were generated by opening the system to Internet Web browsing and by interconnecting different agencies and the development of a security architecture.

The infrastructure in question involves a worldwide corporate network whose primary purpose is to provide connectivity and transport for users. The corporate system does not contain sensitive data itself, but must be available for users. Users maintain their own networks that contain sensitive information. The information is contained on outer networks and inner networks, with the latter containing more critical information typically. National guidance directed that these inner critical networks should be protected at a higher level, as shown in *Table 2*.

A major concern is determining the required protection level. Actuarial models are of little use and the local threat in host countries is of little relevance, given the potential threat from state-sponsored or transnational organisations with worldwide operational capabilities. Protection policies and programmes must be established that will prove adequate for the life-cycle of the system – 10 years or more.

Table 2 shows the solution. The worldwide corporate and sensitive outer networks should be protected to attack level two, whereas the critical inner networks should be protected to attack level four. Essentially, this means that the success of a sophisticated attack against corporate networks or sensitive outer networks is an acceptable risk. This was particularly compelling given the fact that an adversary who is capable of such an attack is unlikely to be motivated to expend resources on such low-value targets.

It is useful to discuss some of the security enhancements that are considered for the protection of more sensitive inner networks. The general approach is referred to as 'high-assurance virtual wide area network' (VWAN) and involves segmenting the architecture and applying several layers of protection. Controls applied at various layers include:

- enhanced identification and authentication;
- application hosting and managed applications;
- restriction of Internet Web browsing;
- restriction of e-mail attachments;
- stringent filtering policies on routers;
- strong personnel, physical and administrative controls; and
- network and host-based monitoring.

Of particular interest is the use of application hosting for the most sensitive systems, providing maximum protection against close access and technical attacks. For example, files do not have to be transmitted over the network and local storage of data can be controlled so that the terminal is non-sensitive when not in use. Physical access controls for personnel and physical control can be used over the location of terminals as an additional level of identification and authentication. In addition, techniques are being developed for identifying and correlating indicators of anomalous activity, whether by individual users or on network segments.

All of these techniques are well understood, although they are implemented for the protection of unclassified systems only rarely because of operational restrictions and inconvenience to users. One of the best reasons for using the design-based approach that is discussed in this article is the fact that it helps to minimise the restrictions users find onerous and, therefore, helps to resolve organisational resistance. ■