Journal of Applied Security Research

ISSN: 1936-1610 (Print) 1936-1629 (Online) Journal homepage: https://www.tandfonline.com/loi/wasr20

A More Rigorous Framework for Security-in-Depth

Rick Nunes-Vaz PhD , Steven Lord PhD & Jolanta Ciuk PhD

To cite this article: Rick Nunes-Vaz PhD , Steven Lord PhD & Jolanta Ciuk PhD (2011) A More Rigorous Framework for Security-in-Depth, Journal of Applied Security Research, 6:3, 372-393, DOI: 10.1080/19361610.2011.580283

To link to this article: https://doi.org/10.1080/19361610.2011.580283

Published online: 15 Jul 2011.

Journal of Applied Security Research, 6:372–393, 2011

ISSN: 1936-1610 print / 1936-1629 online
DOI: 10.1080/19361610.2011.580283

Routledge
Taylor & Francis Group

A More Rigorous Framework for Security-in-Depth

RICK NUNES-VAZ, PhD, STEVEN LORD, PhD, and JOLANTA CIUK, PhD
Counter-Terrorism & Security Technology Centre, Defence Science & Technology Organisation

While the concept of security-in-depth or layered security has a long history, it still lacks clear definition, hampering attempts to identify the most effective target of security enhancement. A rigorous definition of security layer *enables the development of useful principles to guide security investment. Risk minimization is best achieved by strengthening the layer that may already be the most effective, and by focusing on the weakest function within that layer. Moreover, security-in-depth relies not only on generating effective layers, but also on their coherent integration with maintenance, training, protocols and policies, all aligned with management structures and culture.*

KEYWORDS Security-in-depth, security layers, security risk management, terrorism risk

INTRODUCTION AND PURPOSE

Most lay people, and certainly all security professionals, have an intuitive understanding of the phrase *security-in-depth* (or *defense-in-depth*, which we treat synonymously). As a security philosophy, its use is believed to date at least to Roman times (Luttwak, 1976). Security-in-depth usually includes the concept of *layers*, as in *layered security*, although there is only a vague consensus in the literature on the question of what constitutes a layer. These

The authors acknowledge colleagues John Asenstorfer, Wayne Hobbs, Edward Huddy, John Percival, and Debbie Richards, with whom their many conversations on these topics with that have helped to strengthen the concepts and framework.

Address correspondence to Rick Nunes-Vaz, Counter-Terrorism & Security Technology Centre, Defence Science & Technology Organisation, P.O. Box 1500, Edinburgh, South Australia 5111. E-mail: rick.nunes-vaz@dsto.defence.gov.au

ambiguities make it difficult to answer the fundamental question, "From a philosophical, or design principle standpoint, where should investment be targeted in order to create or enhance a security-in-depth system?" Is it wiser to add another layer (which might seem to be the approach adopted in aviation security), strengthen the weakest layer (Schneier, 2003), or perhaps demand that all layers meet certain minimum performance standards?

This article is ambitious in trying to achieve the following multiple objectives. We primarily aim to create a foundation or framework from which the resource allocation question may be addressed, informing strategic investment decisions in security enhancement. The journey, however, requires close examination of concepts associated with security layers and development of a consistent language to support a more rigorous approach.

Background

The simplest conceptual form of security-in-depth is often termed the *security onion* or *bull's eye approach* (Talbot & Jakeman, 2009) and is ubiquitous in at least the physical protection and information technology security domains (e.g., Howard & LeBlanc, 2003; Jaquith, 2007; Viega & McGraw, 2001). For clarity, we restrict attention to physical protection systems, noting that we have not tested the translation of concepts to other security domains; notwithstanding Schneier's (2000) observation that "defense in depth is another universal security principle that applies to computers just as it applies to everything else" (p. 370).

In many strategic and national policy documents (e.g., U.S. Department of Homeland Security, 2007), the meaning of security-in-depth is regarded as assumed knowledge and remains undefined, as in "[t]he Coast Guard will implement a layered defense intended to thwart terrorist threats as far from our shores as possible" (U.S. Coast Guard, 2002, p. 18). More commonly, definitions represent variations on a theme that focuses on redundancy across a suite of security measures (e.g., Robinson, Lake, & Seghetti, 2005), for example, "the [Federal Aviation Administration] set and enforced security rules . . . [that] were supposed to produce a layered system of defense. This meant the failure of any one layer of security would not be fatal, because additional layers would provide backup security" (National Commission on Terrorist Attacks Upon the United States, 2004, p. 83), or "security-in-depth is a multilayered system in which measures combine to support and complement each other, making it difficult for [an attacker]" (Australian Government, Attorney-General's Department, 2007, p. E7). Later in this article, we propose a more robust definition of *security-in-depth*.

Given that the concept of *security-in-depth* is almost universally seen as synonymous with *layered defense*, it implies that depth is created by multiple (complementary and partially redundant) layers. However, despite this consensus position, the term *layer* remains undefined in the literature, which makes the task of measuring layer (and security-in-depth) effectiveness

difficult. Nevertheless, the literature contains many different examples or models of layered systems that are considered to satisfy the inferred requirements of security-in-depth. For example, national security strategies usually use implied complementary layers, using terminology such as "prevent," "prepare," "respond," and "recover" (Australian Government, 2005; Australian Government, Attorney-General's Department, 2011), "prevent, pursue, protect, and prepare" (British Foreign & Commonwealth Office, 2006, p. 1), or "prevent and disrupt, protect, respond and recover" (U.S. Department of Homeland Security, 2007, p. 1). Also proposed as layers are variants of the concepts "people," "procedural" and "technical" (Blackwell, 2008; Trusted Information Sharing Network, 2008), or "physical," "procedural" and "psychological" (Royal Canadian Mounted Police, 2004) and similar constructs. Alternatively, there are variations on themes built around functions such as "deter," "deny," "delay," "detect," "respond" and "recover" (e.g., Australian Government, Department of the Prime Minister and Cabinet, 2010; Dillon, Liebe, & Bestafka, 2009; EURIM, 2010; Norman, 2010; Royal Canadian Mounted Police, 2004; Talbot & Jakeman, 2009; Trusted Information Sharing Network, 2008; U.S. Department of Homeland Security, 2009), "detect, assess, warn and defend" (U.S. Department of Defense, 2003, p. 8) or "surveillance, reconnaissance, tracking and interdiction" (U.S. Coast Guard, 2002, p. 18).

As the layers or functions such as deter, prevent, and respond may require the implementation of several complementary (e.g., physical, procedural, technical) elements, they should be seen as complex concepts in themselves. They should consequently be distinguished from security controls (also security measures) that represent specific elements or security solutions such as contraband screening or identity management systems. However, the distinction between *layer* and *control* is frequently lost, and as a consequence, so too is the understanding of effective security-in-depth, as seen in "... a layered approach to counter-terrorism is essential because no single measure will be fully effective ..." (Australian Government, Department of the Prime Minister and Cabinet, 2010, p. 19), "... no single security measure is foolproof. Accordingly, the [Transportation Security Administration] must have multiple layers of security ..." (National Commission on Terrorist Attacks Upon the United States, 2004, p. 392), or "the Government of Canada added another layer of security to the nation's aviation system by unveiling a new program to screen non-passengers ..." (Office of the Auditor General of Canada, 2004, p. 47).

Suggested Terminology

Security-in-Depth

Security-in-depth (or defense-in-depth) is a strategic concept. The heart of security-in-depth involves the design and coordinated implementation of multiple security controls into layers, to ensure that an attack of any type[1]

cannot succeed (or an accident cannot proceed) without defeating all security layers. This goes further than Garcia (2006, p. 40), who said "protection-in-depth means that, to accomplish the goal, an adversary should be required to avoid or defeat a number of protective devices in sequence."

Security Layer

This article argues for a clear distinction between *security layer* and *security control* (device or system) on the grounds that a collection of controls must satisfy additional properties to warrant the term *layer*, because security-in-depth relies on the attacker defeating layers not controls. The case is developed in the body of this article that a security layer is a set of controls that can potentially stop a defined event from occurring or can entirely eliminate its harmful consequences. Thus, a bomb-detection system can neither defeat the threat and prevent detonation nor mitigate its consequences should detonation occur, which means it is a control, not a layer.

This definition is consistent with the definition of an independent protection layer (International Organization for Standardization & International Electrotechnical Commission, 2009), as a "device system or action that is capable of preventing a scenario proceeding to its undesired consequence," which has its origins in hazard analysis associated with the process industry (International Organization for Standardization & International Electrotechnical Commission, 2003, p. 60). It implies that any individual security layer alone, by either reducing the likelihood to zero or reducing the consequence to zero, may eliminate the risk. In general, security layers are imperfect and hence complete elimination of risk will rarely be achievable. Nevertheless, a layer possesses all of the attributes to make this result theoretically possible. Placing such a stringent constraint on the use of the term layer creates a more rigorous foundation for assessing security system effectiveness while also adding clarity to the security lexicon,[2] but this distinction does not appear to exist in the security literature to date.

Security Function

Layers require the coordination and combined effects of security functions. Functions in themselves can neither defeat threats nor eliminate their consequences, but they should be seen as generic components of security layers and include, for example, deter, detect, delay, alert, neutralize, contain, and restore (Garcia, 2006, 2008).

Security Control

Also termed *countermeasure* or *barrier* (used synonymously here), a *security control* is a physical, psychological, procedural, technical, or other device that performs or contributes to one or more security functions.

Risk Management

Another key consideration in developing or enhancing a security system hinges on the strategic objectives. The rhetoric of federal agencies in many Western nations has shifted substantially in the last decade toward a risk minimization stance (e.g., Cabinet Office, 2010; Chertoff, 2005). This is a useful perspective because it considers security throughout its timeline from emergence of the threat, through the occurrence of a defined security event, to its effects and downstream consequences—as represented in the bow-tie diagram (see Figure 1).

Controls may intervene between the threat and a prescribed event,[3] or they may support management of effects and downstream consequences. The path of attack shown in Figure 1 is intended to indicate that a particular event may be one manifestation of a range of possible threats and attack pathways, and the nature of consequences depends on the timeliness and effectiveness of incident response and consequence management controls. The risk, therefore, depends on the effectiveness of the system of controls as a whole.

Formalizing the effects of security controls into a single vulnerability parameter (as in Dillon et al., 2009; Garcia, 2006; Norman, 2010; Willis, Morral, Kelly, & Medby, 2005), expresses the notion that enhanced security reduces risk by reducing vulnerability, as follows:

$$Vulnerability = f\,(security\ controls)$$

and

$$Risk = Likelihood \times Consequence \times Vulnerability$$

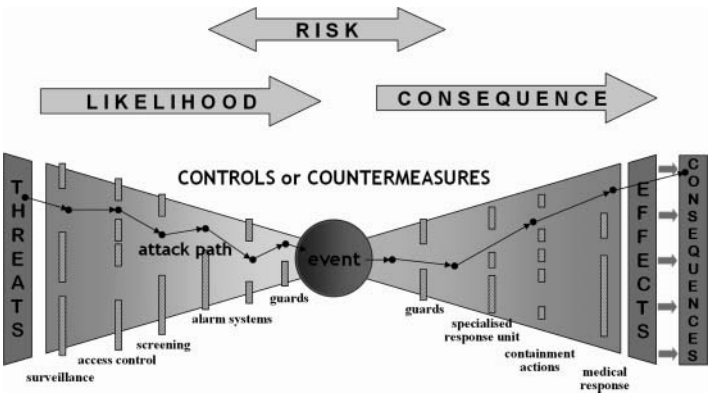

FIGURE 1 The bow-tie diagram conceptually separates the role of controls that affect the likelihood of a security event, from those that manage its consequences. Each control acts (in an imperfect manner as illustrated by broken barriers impeding the left-to-right progress of the event's timeline) as a contribution to overall risk reduction (adapted from Talbot & Jakeman, 2009).

in which likelihood and consequence represent the threat and context and are independent of changes to security control implementation. In practice, vulnerability is usually considered to influence only the likelihood of attack, that is, security controls reduce likelihood either by reducing the attractiveness of the target or by making a successful attack harder to achieve. However, there is another component of vulnerability that tends not to be considered. It is the susceptibility of the target once an attack has occurred, and the extent to which the event may trigger a cascade of harmful consequences. Security controls may be used to mitigate consequence, for example, in the use of physical shields to protect fixed assets from explosive blast damage, or cordons and response protocols to contain any further harm. This second component of vulnerability is to some extent independent of the former, and it tends to affect consequence rather than likelihood. This view is more closely aligned with the Australian/New Zealand Risk Management Standard's Handbook (Australia-HB 167, 2006), in which controls contribute directly and independently to likelihood and consequence (and thereby risk) reduction, as in Figure 2, precluding the need for explicit representation of vulnerability.

Thus it is more useful to think in the following terms:

$$Likelihood = f_L\ (security\ controls)$$
$$Consequence = f_C\ (security\ controls)$$

Rather than using a simple product, risk is more generally considered to be a function of likelihood and consequence:

$$Risk = f_R\ \{Likelihood,\ Consequence\}$$

and thereby

$$Risk\ (security\ controls) = f_R\ \{f_L\ (security\ controls),\ f_C\ (security\ controls)\}$$

In the previous formulas, the functions represent not only the presence of controls, but also factors associated with their performance and effectiveness

FIGURE 2 Security controls influence both likelihood and consequence independently.

resulting from organizational attributes relating to maintenance, training, clarity of roles and responsibilities, and protocols and policies. These broader aspects of security implementation and the creation of security depth are discussed in the "Dimension of Security-in-Depth" section, after first establishing a framework that formally situates controls within layers.

MEASURING THE EFFECTIVENESS OF A SECURITY-IN-DEPTH SYSTEM

It is common, particularly associated with the security of information technology and safety critical systems such as nuclear facilities (Australian Government, Attorney-General's Department, 2007; Blackwell, 2008; Royal Canadian Mounted Police, 2004), to see controls classified in terms of their contributions to managing physical, personnel, procedural, and technical security, as in, "the concept of defence in depth ... consists of implementing several layers of defence, including both administrative aspects ... and technical aspects ..." (International Atomic Energy Agency, 2009, p. 8).

These properties or attributes of the security system are not layers, and the primary issue with a taxonomy of this type lies in assessing the effectiveness of the system. The question revolves around defining an appropriate balance or profile of contributions across the components. Because these are not measurable properties (other than by aggregating ticks associated with contributions from controls), it is difficult to meaningfully compare the effectiveness of two different profiles. Nor is there any guidance on what constitutes an effective profile for any particular context and threat.

A more meaningful approach (Australian Government, Attorney-General's Department, 2010; Ayyub, McGill, & Kaminskiy, 2007; Dillon et al., 2009; Garcia, 2008; Norman, 2010) weighs contributions from security functions, such as deter, detect, alert, and respond. In this case, an aggregate assessment of the effectiveness of each function leads to a meaningful measure of probability (see Figure 3). For example, one or more detection devices will provide a probability of detection for a defined threat (as assessed in an operational context), noting that these probabilities will change with the threat, as a control may be highly effective for one threat but relatively ineffective for another.

Measurement typically involves aggregation within columns to assess the adequacy of each function, and whether a balance has been achieved across all required functions (e.g., Norman, 2010, p. 336). This does not usually extend as far as assessing scores in relation to each threat individually because effectiveness is commonly assessed using a subjective (e.g., three- or four-point) scale and the additional accuracy may not be justified.

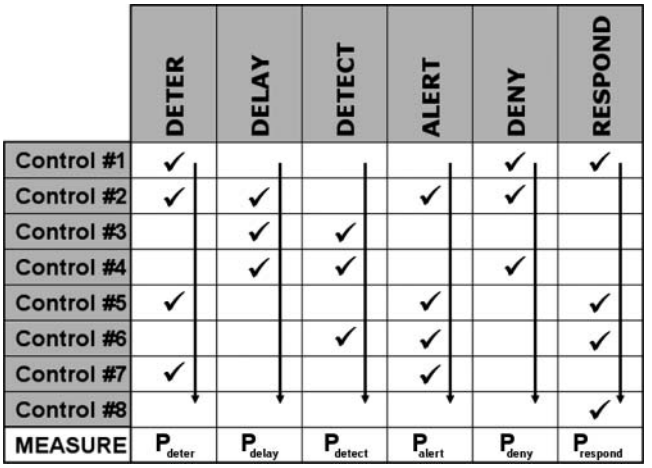

FIGURE 3 A way to represent the contributions of security controls to the performance of security functions.

However, again noting the meaning of security layer, all such approaches suffer the limitation that even perfect detection (or delay, or alert, or response) cannot alone prevent a security incident or mitigate its effects. Thus aggregating contributions to these functions within isolated columns fails to represent the effectiveness of security. An effective neutralization response requires a timely alert following initial detection, and may also demand a degree of delay imposed on the attacker once detection has occurred. Thus, the probability of preventing a scenario from proceeding to its undesired consequence should be assessed as the cumulative or joint probability associated with the interactions of several functions, and their supporting controls.

Accordingly, these sorts of taxonomies or classification approaches, which exhaust the scope of the current literature, provide no guidance on how to enhance layered security.

Dependent and Independent Layers and Functions

To clarify the concept of security layers, consider for the moment only those layers that affect likelihood. Let L_{prior} denote the chance of a threat engaging with n independent (serial) layers of security. Let P_i denote the effectiveness of the ith layer given engagement, where effectiveness is the chance of stopping the event. The contribution of a series of such layers may be represented in the following form:

$$L = L_{prior} \times \prod_{i=1}^{n} (1 - P_i) \tag{1}$$

where Π represents a product, and L is the chance of the event. This representation, which is consistent with Garcia (2008) and others, formalizes the definition of a layer such that a single fully effective layer ($P_i = 1$) eliminates risk by reducing likelihood to zero, regardless of the contributions from other layers.

As we have seen, security functions (e.g., detect) do not generally constitute layers in themselves, but they contribute to layer effectiveness. If a layer relies on the effects of a series of dependent functions, where the effectiveness of the *jth* security function (i.e., the chance of detecting, alerting, responding) is Q_j the simplest representation of the effectiveness, P, of the layer is the following:

$$P = \prod_{i=1}^{m} Q_j. \tag{2}$$

This implies that a complex layer involving several (m) interdependent functions demands much greater effectiveness in each of its components to achieve the same overall effectiveness as a simple layer involving few functions.

By way of an example, consider a layer that relies on three functions: this might be a PREVENT layer that requires coordinated contributions from detection, alert and neutralization. If each function's effectiveness is 50%, then the layer has an overall effectiveness of (0.5 × 0.5 × 0.5) or 12.5%, and prior likelihood (in this one-layer system) is modestly reduced by the same 12.5%. More commonly, the literature considers security functions as independent layers in themselves. If detection, alert, and neutralization (each with the same 50% effectiveness) were treated as layers, then Equation (1) implies that prior likelihood is reduced by 87.5% and thus the risk is almost eliminated. The distinction between layer and function is therefore critical to valid appraisal of system effectiveness. This approach accords with intuition in that a layer that relies on many functions to be effective, requires each of those functions to be effective. If each function has the same effectiveness (e.g., 50%) then the greater the number of connected functions ("moving parts"), the greater the risk.

Security-in-Depth Design Principles

This formalization of security-in-depth and its representation of security layers indicates that security risk may be minimized (with respect to any single threat) by maximizing the effectiveness of any one security layer, rather than balancing the contributions of several layers. This is counter-intuitive with regard to the folklore surrounding security-in-depth. There may well be several security functions required to create the layer, underpinned by a multitude of security controls, but investment may be best targeted at the

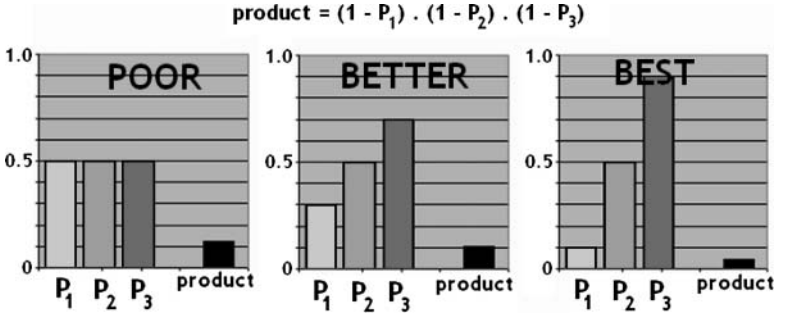

FIGURE 4 To minimize the product (or risk) as the outcome of several security layers—according to Equation (1) in the text—the best strategy is to maximize the effectiveness of any single layer.

layer which is already the most effective (see Figure 4), rather than spreading the investment across all layers.

A need to balance investment across many layers does, however, normally arise. First, because it is difficult to build a perfect layer and investment in strengthening an effective layer may suffer from a diminishing returns effect. Second, because different layers are more or less effective against different threats, the wider the spectrum of threats under consideration (which we call *scenario depth* in the "Dimensions of Security-in-Depth" section), the broader the investment strategy will generally need to be. Nevertheless, the rationale for investment should remain traceable with regard to the range of concerns.

Within any single security layer, Equation (2) states that layer effectiveness is maximized by maximizing the product of the effectiveness of its internal functions. In this case, the product is maximized by balancing across all functions (Figure 5), that is, the layer is most seriously compromised by weakness in any function. Investment to enhance security should therefore be focused on eliminating weakness (in accord with Schneier, 2000) in the internal functions of an important security layer.

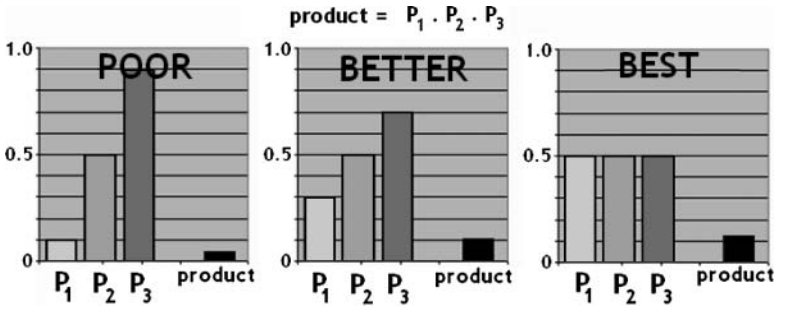

FIGURE 5 To maximize the product (or effectiveness) of several security functions— according to Equation (2) in the text—the best strategy is to eliminate weakness in any single function.

Alternatively, risk reduction may be achieved by simplifying the chain of dependencies between interdependent security functions.

SECURITY LAYERS TO MANAGE THE RISK OF TERRORISM

The Likelihood of Attack

Consider the layers and functions required to stop a terrorist attack, as represented in Figure 6. It is apparent that the DETER function is, in fact, a layer because it can (and does) stop potential attacks independently of other functions (e.g., it does not require detection, alert, or neutralization, other than through the attacker's perception). Note, however, that many authors caution against attempts to measure and rely on deterrence (e.g., Norman, 2010). Detection, alert, neutralization and perhaps delay are all required to provide effective prevention for those attackers who are not deterred.

Figure 6 illustrates the roles of these security functions. Thus, if attacked, failure to detect leads to a harmful event. If detected, failure to alert leads to the event and, if alerted, failure to neutralize leads to the event.

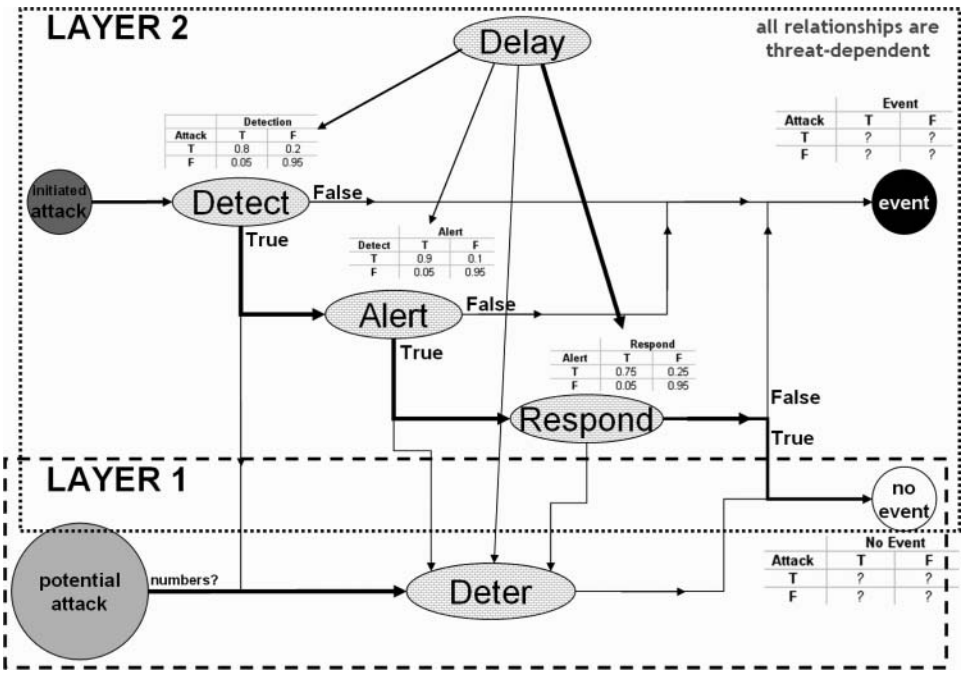

FIGURE 6 An illustration of the manner in which security controls, through their contributions to security functions, may intervene to prevent the occurrence of an attack (shown as the "event"). As discussed in the text, the event may be stopped by deterrence alone (LAYER 1), but if not deterred, it requires the combined effects of all other functions to achieve prevention (LAYER 2). Performance tables associated with functions and outcomes are purely illustrative.

The performance characteristics of each function (derived from the combined effects of several controls) should be evaluated in the appropriate operational contexts, and may be represented by a table that shows results for each of four possible outcomes. For example, the detect function's four outcomes are (a) detection given attack, (b) detection given no attack (or false positive), (c) no detection given attack (or false negative), and (d) no detection given no attack. The probability of an alert given a detection may be similarly assessed, and so too the probability of neutralization given an alert. Hypothetical outcomes of these assessments are included as tables in the figure.

With this construct, it is clear that the delay function contributes to security indirectly by (potentially) changing the performance of each of the core functions. For example, delaying an attacker after initial detection may provide sufficient time to mount, and therefore increase the probability of achieving, an effective neutralization. Thus, detection, alert, neutralization, and delay all contribute to a single security layer, which we call the PREVENT layer.

The analysis indicates that in managing the risk (to this point considering only the likelihood) that a defined security incident will occur, security controls contribute to two relevant security layers: deterrence and prevention. Prevention is itself made up of several functions that (depending on the context and the threat) may include detection, alert, neutralization, and delay, but these should not be considered an exhaustive set.

These principles were recognized in the design of security systems intended to prevent attack at Sandia National Laboratories, which is responsible for ensuring the safety of the U.S. nuclear arsenal: "a physical protection system must accomplish its objectives by either deterrence or a combination of detection, delay and response" (Garcia, 2006, p. 35).

The Consequences of Attack

A similar postevent analysis reveals a further two relevant security layers, which we term *PROTECT* and *CONTAIN*.

Protection manages (and potentially eliminates) the consequences of attack by putting controls in place prior to the event. For example, vaccination of a vulnerable population may render a specific biological attack ineffective. Similarly, the use of barriers (to create standoff), blast shields and protective films can provide protection of fixed assets against explosive devices. Protection, as used here, requires no active intervention when an event occurs, and is useful for events where no warning is anticipated.

The CONTAIN layer is a combination of incident response and consequence management capabilities and actions, represented by additional detect, alert, limit, and respond functions, although other functions may

also be relevant depending on the threat. The time horizon applied here excludes the slower cascading impacts (e.g., loss of business revenue because of behavioral changes induced by fear) and recovery activities (e.g., rebuilding physical facilities) that play out in the days to months that follow a significant incident. Thus recovery or restoration is considered not to be a component of security as such.

The Four Layers Required to Manage Terrorism Risk

This discussion leads to a representation (Figure 7) of the whole security system in terms of four layers: DETER, PREVENT, PROTECT, and CONTAIN. Two layers manage pre-event likelihood, and two manage post-event consequences.

Note that DETER and PROTECT are seen as largely passive layers; they require no actions following the initiation of an attack. They are also not explicitly divided into example functions. While many factors and functions (e.g., delay, in the form of walls and razor wire) contribute to deterrence, the perceived effectiveness of the remaining security functions and layers generates deterrence, and those elements are already represented in other layers.

In the case of PROTECT, the nature of protection is highly dependent on the context and the threat, and cannot be specified in greater detail while remaining generic.

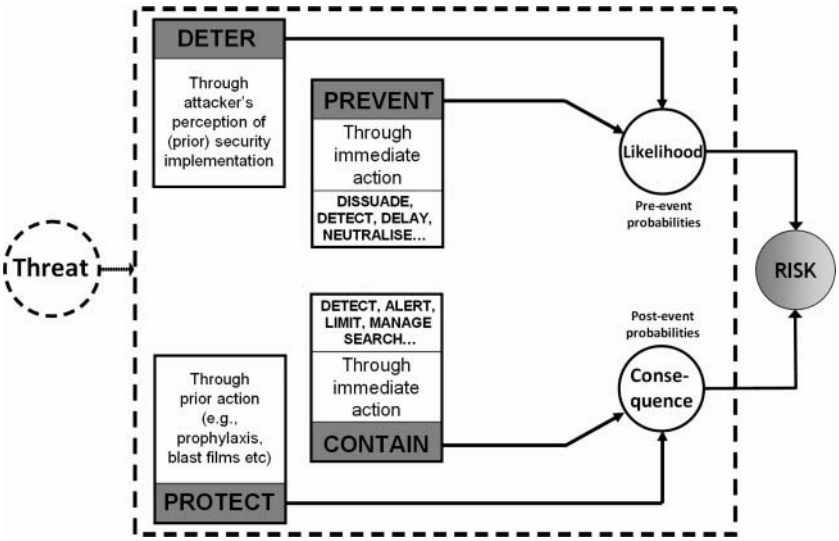

FIGURE 7 Representation of security layers relevant to terrorism, as discussed in the text.

The layers and functions are implemented by controls that are subject to many complex interdependencies. These connections are believed to sit primarily within single layers, but may stretch across layers, for example, pre-event detection may assist post-event containment. Note that PREVENT (see Figure 7) contains a function called *dissuade*, which is related to deterrence. It acknowledges that active prevention (e.g., lights turning on to indicate detection, an intruder alarm sounding, or the appearance of guards) may prompt a perpetrator to abort an attack in progress.

Security Risk Design Principles

From a design perspective, if the objective is to minimize the likelihood of a specific attack, then investment should be directed toward maximizing the effectiveness of either the DETER or PREVENT layers. Because "... one cannot estimate deterrence from the existence of countermeasures [controls] ..." (Norman, 2010, p. 333) and deterrence relies on the perception of security that arises from evident strength in other layers, the focus of investment should be directed towards PREVENT. This implies that all functions of prevention should be examined, and investment should focus on raising the effectiveness of the weaker components in the system (see the "Design Principles" section).

If it is acknowledged that particular types of attack are difficult to prevent, then the focus should shift to maximizing either the PROTECT or CONTAIN layers. Where protection is feasible it may be preferable because, once in place, it does not rely on the simultaneous coordinated actions of multiple players (or agencies) and, for this reason, may be more reliable. However, protection may be difficult to achieve, for example, protecting personnel in the open from the effects of an explosive device.

It is apparent that different threats imply differing strategies with respect to targeting particular functions and layers for enhancement. Furthermore, each control will generally contribute to more than one function and perhaps more than one layer. Investment strategies must therefore account for (a) the need to target specific layers according to the threat spectrum under consideration, (b) the contributions of each new or enhanced control to all relevant layers, and (c) the benefit-cost of each potential investment in the overall security of the system. In this way, investment analysis requires in-depth assessment of the various alternative portfolios of control options.

It is also relevant to note the implications of differences in stance (strategy) adopted by national security decision-makers in government, compared with those who own or manage assets such as critical infrastructure. Government officials must manage risk with a great deal of uncertainty about the targets of potential attack. They are therefore hamstrung by resource constraints and cannot invest in target-specific prevention or protection of

vulnerabilities. They must focus on pre-event risk management, specifically deterrence and detection (intelligence assessments of the magnitude of the threat and, if possible, the intended target). Infrastructure owners, however, know precisely what the potential target is and can therefore invest in PREVENT, PROTECT, and CONTAIN layers according to external (government) assessments of the threat magnitude. These stakeholders create security-in-depth through their cooperation and shared management approach.

DIMENSIONS OF SECURITY-IN-DEPTH

Previous sections have presented a case for the adoption of a consistent language to describe security systems, and a framework within which to assess the relative values of alternative security enhancement packages. However, effective security-in-depth demands much more than the implementation of security controls within defined layers. Their effectiveness depends intimately on the coherence of the systems with regimes of maintenance, training, protocols, and policies, which must all be aligned with management structures and culture. Thus, depth of security may be characterized by depth or robustness in three conceptual domains: *system depth*, which relates to the configuration of the control systems themselves; *scenario depth*, which characterizes the spectrum of threats that the systems are designed to deal with; and *structural depth*, which represents all of the additional organizational considerations required to ensure that the systems perform at their optimal or acceptable levels.

System Depth

In analogy with the technical view of safety instrumented systems (International Organization for Standardization & International Electrotechnical Commission, 2003), the use of diverse controls to provide security may be seen as a form of system integration problem—involving human and technical systems. At the most fundamental level (Fleming & Silady, 2002), independence is required to ensure that separate controls do not share a single cause of failure, such as the loss of electrical power or loss of communications. Regard for human factors associated with operators' abilities to discriminate intent, behaviors, and different objects as well as their cognitive endurance in these roles, are just some of the issues that go to the heart of system performance.

The concept of system depth (see Figure 8) as a primary component of security-in-depth, is intended to convey the notion that the control systems themselves, as well as meeting the needs of security, must be integrated to ensure systemwide robustness and resilience in their operational context.

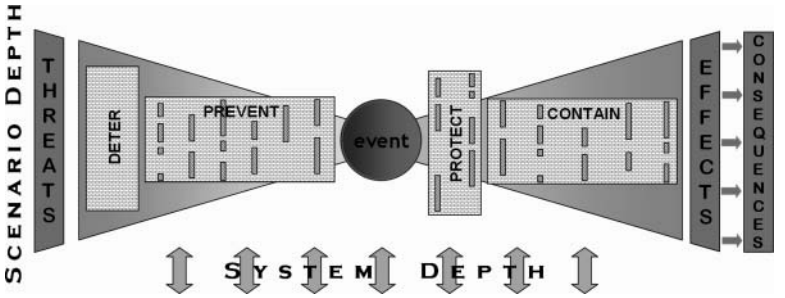

FIGURE 8 A summary view of concepts discussed in the text, showing that security controls, relating to specific terrorism scenarios discussed, contribute to four pre- and postevent security layers.

Randomization of the timing, pathways, and even numbers of guards, for example, enhances the system's robustness, and generates system depth, particularly if closed-circuit television operators are primed to observe complementary zones with regard to patrols.

Scenario Depth

Scenario depth is intended to ensure that the systems have been designed and tested to deal with the appropriate range of scenarios. In general, it is insufficient to consider a broad type of threat, such as vehicle-based improvised explosive device, without considering the scope of alternative delivery modes and targets and how each may play out within the context in question. Scenario analysis should therefore go hand-in-hand with the assessment of existing security systems, and any considerations of their enhancement.

Structural Depth

Following the concepts of Fleming and Silady (2002) and Murphy and Pate-Cornell (1996), the third dimension of security-in-depth is here termed structural depth (see Figure 9). The base of the structure, called the system or S level, contains the coalface control systems that perform the functions of detection, alert, and neutralization, and should be considered to include controls such as guards.

Immediately above the system level is the action, or A level, which represents the decisions and actions of individuals as they affect performance at the S level. They include operational procedures and protocols, and all training, testing and maintenance aspects of the system. The A level enables the S level by ensuring that systems are always operational (or backups are

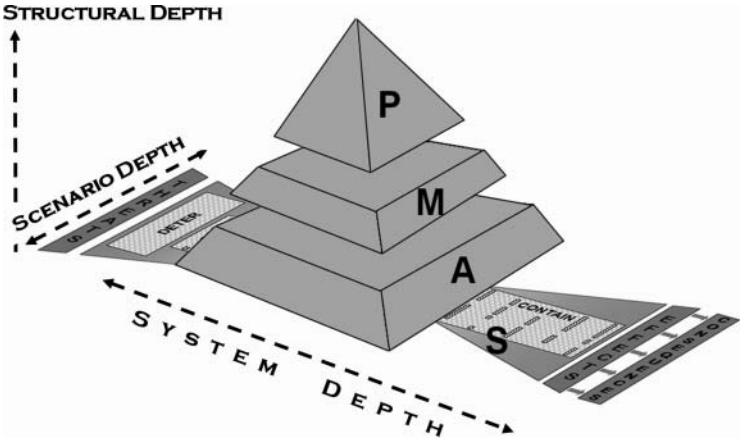

FIGURE 9 The System-Action-Management-Policy (or S-A-M-P) structure that contributes to security-in-depth. The text argues that security-in-depth arises from coherence among system, scenario, and structural depth through the effective implementation of measures at all levels.

integrated), and operators are well trained and have clear guidance on all aspects of their roles and responsibilities.

Above the action level is the management or M level (Murphy & Pate-Cornell, 1996), which ensures the following: (a) that the organization is appropriately structured to achieve its objectives; (b) that command and control is well defined through all possible states of the system; (c) that strategic, operational, and emergency response plans are developed and in place; (d) that the system is appropriately tested and evaluated; and (e) that capability development is effectively managed and resourced.

Last, at the top of the structure is the policy or P level (not part of the Murphy & Pate-Cornell, 1996, scheme), which sets the strategic parameters that define the system's objectives and culture, and which drives consistent investment and development throughout the structure. It follows that effective security-in-depth requires coherent implementation of controls and measures throughout the S-A-M-P structure.

Note also the vertical connectivity implied in the three-dimensional structure. For example, a change of policy or objectives at the P level (perhaps because of the inclusion of a new threat) implies adjustment of management, plans and resources (at M), requiring new training, testing and decision protocols (A) and potentially new systems (S). The structure is appropriately driven top-down, but may require reappraisal bottom-up if new systems are acquired, which may occur when systems are replaced at the end of their operational lives by others that operate according to different principles.

SUMMARY

Despite its long history in military conflict, and its ubiquitous usage in the security community, the concept of security-in-depth is at best ambiguously described in the literature. This article develops a structured framework for describing the relations between security devices (controls) and the layered security (security-in-depth) system of which they form a part in order to inform resource allocation decisions on security enhancement.

In this framework (which is consistent with risk management and security risk management standards), security controls are physical, psychological, procedural, technical, or other devices that perform or contribute to one or more security functions. Security *functions*, such as detect, delay, alert, and respond, are combined in a coordinated (designed) manner to create security layers. Each security *layer* is a set of controls performing functions that, when effectively integrated and enabled, can potentially eliminate risk either by stopping a defined event from occurring or precluding its harmful consequences. Thus, *security-in-depth*, as defined here, is the designed and coordinated implementation of multiple security controls into layers, to ensure that an attack of any type cannot succeed without defeating all security layers. A simple analysis founded on this framework is used to show that common misconceptions about the distinctions between security functions and security layers can lead to substantial overestimation of security effectiveness.

The principles of layered security implied by the framework are illustrated by considering the issue of securing a defined facility such as a military base or an infrastructure facility from the threat of terrorist attack by improvised explosive device. For this case, it is shown that security controls may be configured into four relevant layers, termed DETER, PREVENT, PROTECT, and CONTAIN. Two of the layers (DETER and PREVENT) target likelihood reduction and focus on stopping the event. If deterrence fails (as an independent layer) then prevention is intended to halt the attack's progress. The other two layers (PROTECT and CONTAIN) focus on consequence management. These layers make distinctions between risk treatments performed in advance of an attack, such as the installation of blast shields and protective films on fixed assets, from those which require active intervention during an attack (alerting vulnerable personnel, medical interventions) to contain the harm. If (prior) protection fails, then containment is intended to limit harmful consequences. Thus, risk may be eliminated by any single fully effective layer although, in practice, achieving 100% effectiveness in any layer may be difficult or costly. Other threat types and security contexts may warrant alternative security-in-depth (layer) constructs.

The analysis developed here implies guiding principles for investment to enhance security. It is shown that, contrary to intuition, risk reduction may be maximized by focusing investment within a single layer relevant

to a specific threat, even though this may already be the most effective layer. Once the optimum layer has been identified, its effectiveness may be maximized by investing to strengthen the weakest security function within the layer. As different threats imply the need to strengthen different layers (e.g., prevention may be preferred for terrorism, while containment may be the more effective choice for dealing with a rogue insider), investment strategies must account for (a) the need to target specific layers according to the threat spectrum under consideration, (b) the contributions of each new or enhanced control to all relevant layers, and (c) the benefit-cost of each potential investment in the overall security of the system. In this way, investment analysis requires in-depth assessment of the various alternative portfolios of control options.

Having established a framework for organizing security controls within layers, it is argued that the effective implementation of security-in-depth demands a much more holistic approach to enabling all aspects of the security system to perform as required. This broader interpretation, encompassing technical, procedural, managerial, organizational, and cultural factors, is represented through a conceptual hierarchy called the system-action-management-policy (SAMP) structure (Murphy & Pate-Cornell, 1996). The system level (S level) contains the security controls and layers as discussed above. The action level (A level) represents the decisions and actions of individuals as they affect performance at the S level. They include operational procedures and protocols, and all training, testing, and maintenance aspects that ensure the systems are always operational, and operators are effective. Above the A level is the management level (M level), which ensures the following: (a) that the organization is appropriately structured to achieve its objectives; (b) that command and control is well defined through all possible states of the system; (c) that strategic, operational, and emergency response plans are developed and in place; (d) that the system is appropriately tested and evaluated; and (e) that capability development is effectively managed and resourced. Last, at the top of the structure is the policy level (P level), which sets the strategic parameters that define the system's objectives and culture, and which drives consistent investment and development throughout the structure. Security-in-depth is achieved through coherent implementation of sociotechnical systems at all levels.

It is interesting to reflect that the national security strategies of various nations are described using an implied layered construct such as "prevent, prepare, respond, recover" (Australian Government, Attorney-General's Department, 2011), or "prevent, pursue, protect, prepare" (British Foreign & Commonwealth Office, 2006). The implementation of national security according to these constructs is considered to be "comprehensive" (Australian Government, Department of the Prime Minister and Cabinet, 2006, p. 3), although the functional composition of each layer appears not to be articulated in more detail than defining the scope of threats and the agencies

involved. It would be an interesting exercise, although beyond the scope of this article, to overlay their implementation onto a formal security-in-depth framework as defined here to assess for completeness and effectiveness.

NOTES

1. Within a defined scope.
2. To make this distinction more apparent in the text, layers appear in capital letters.
3. It is also important to be precise about defining the event itself.

REFERENCES

Australian Government. (2005). *National Counter-Terrorism Plan*. Canberra, Australia: Author. http://www.ag.gov.au/agd/WWW/rwpattach.nsf/VAP/(5738DF09EBC4B7EAE52BF217B46ED3DA)~NCTP_Sept_2005.pdf/$file/NCTP_Sept_2005.pdf

Australian Government, Attorney-General's Department. (2007). *Protective Security Manual*. Canberra, Australia: Author. Retrieved from http://www.ag.gov.au/www/agd/agd.nsf/Page/Nationalsecurity_ProtectiveSecurityManual (PSM2005)

Australian Government: Attorney-General's Department. (2010). *Protective Security Policy Framework*. Canberra, Australia: Author. Retrieved from http://www.ag.gov.au/www/agd/agd.nsf/Page/Nationalsecurity_ProtectiveSecurityManual (PSM2005)

Australian Government: Attorney-General's Department. (2011). *Emergency management approaches*. Canberra, Australia: Author. Retrieved from http://www.ema.gov.au/www/emaweb/emaweb.nsf/Page/EmergencyManagement_EmergencyManagementApproaches

Australian Government, Department of the Prime Minister and Cabinet. (2006). *Protecting Australia against terrorism, Australia's National Counter-Terrorism Policy & Arrangements*. Canberra, Australia: Author. Retrieved from http://www.australianislamistmonitor.org/uploads/docs/paat_2006.pdf

Australian Government, Department of the Prime Minister and Cabinet. (2010). *Securing Australia, protecting our community (Counter-Terrorism White Paper)*. Canberra, Australia: Author. Retrieved from http://www.dpmc.gov.au/publications/counter_terrorism/docs/counter-terrorism_white_paper.pdf

Australia-HB167. (2006). Standards Australia/Standards New Zealand. *Handbook: Security Risk Management*. Sydney, NSW, Australia.

Ayyub, B. M., McGill, W. L., & Kaminskiy, M. (2007). Critical asset and portfolio analysis: An all-hazards framework. *Journal of Risk Analysis*, *27*, 789–801. DOI: 10.1111/j.1539-6924.2007.00911.x

Blackwell, C. (2008). *A multi-layered security architecture for modelling critical infrastructure, 2008*. Proceedings of the Seventh European Conference on Information Warfare, June 2008, University of Plymouth, Plymouth, UK, pp. 17–24.

British Foreign & Commonwealth Office. (2006). *Countering international terrorism: The United Kingdom's strategy*. Norwich, UK: Author. Retrieved from http://www.fco.gov.uk/resources/en/pdf/contest-report

Cabinet Office. (2010). *HMG security policy framework* v5.0. London, UK: Cabinet office. Retrieved from http://www.cabinetoffice.gov.UK/sites/default/files/resources/hmg-security-policy_O_O.pdf

Cabinet Office. (2010). *The National Security Strategy of the United Kingdom: A strong Britain in an age of uncertainty*. Norwich, UK: Author. Retrieved from http://www.cabinetoffice.gov.uk/sites/default/files/resources/national-security-strategy.pdf

Chertoff, M. (2005). *Press Release: Homeland Security Secretary Michael Chertoff announces six-point agenda for Department of Homeland Security*. U.S. Department of Homeland Security. Retrieved from http://www.dhs.gov/xnews/releases/press_release_0703.shtm

Dillon, R. L., Liebe, R. M., & Bestafka, T. (2009). Risk-based decision making for terrorism applications. *Journal of Risk Analysis*, *29*, 321–335. DOI: 10.1111/j.1539-6924.2008.01196.x

Fleming, K. N., & Silady, F. A. (2002). A risk-informed defense-in-depth framework for existing and advanced reactors. *Reliability Engineering and System Safety*, *78*, 205–225. DOI: 10.1016/S0951-8320(02)00153-9

Garcia, M. L. (2006). *Vulnerability assessment of physical protection systems*. Burlington, MA: Elsevier Butterworth-Heinemann.

Garcia, M. L. (2008). *The design & evaluation of physical protection systems*. Burlington, MA: Burlington, MA: Elsevier Butterworth-Heinemann.

Howard, M., & LeBlanc, D. (2003). *Writing secure code*. Redmond, WA: Microsoft Press.

International Atomic Energy Agency. (2009). *Security in the transport of radioactive material* (International Atomic Energy Agency Nuclear Security Series No. 9 Implementing Guide). IAEA, Vienna, Austria: Author. Retrieved from http://www-pub.iaea.org/MTCD/publications/PDF/pub1348_web.pdf

International Organization for Standardization & International Electrotechnical Commission. (2003). *Functional safety—Safety instrumented systems for the process industry sector*. IEC 61511-3 Ed1.0, International Electrotechnical Commission, Geneva, Switzerland.

International Organization for Standardization & International Electrotechnical Commission. (2009). *Risk management—Risk assessment techniques*. IEC/ISO 31010: 2009, International Electrotechnical Commission, Geneva, Switzerland.

Jaquith, A. (2007). *Security metrics: Replacing fear, uncertainty and doubt*. Upper Saddle River, NJ: Addison-Wesley.

Luttwak, E. (1976). *Grand strategy of the Roman Empire: From the first century A.D. to the third*. Baltimore, MD: The Johns Hopkins University Press.

Murphy, D. M., & Pate-Cornell, M. E. (1996). The SAM framework: Modeling the effects of management factors on human behaviour in risk analysis. *Journal of Risk Analysis*, *16*, 501–515. DOI: 10.1111/j.1539-6924.1996.tb01096.x

National Commission on Terrorist Attacks Upon the United States. (2004). *The 9/11 Commission Report*. New York, NY: W.W. Norton & Company. Retrieved from http://govinfo.library.unt.edu/911/report/index.htm

Norman, T. L. (2010). *Risk analysis and security countermeasure selection*. Boca Raton, FL: CRC Press.

Office of the Auditor General of Canada. (2004). *Report of the Auditor General of Canada to the House of Commons*. Ottawa, ON, Canada: Author. Retrieved from http://www.oag-bvg.gc.ca/internet/docs/20040303ce.pdf

Robinson, W. H., Lake, J. E., & Seghetti, L. M. (2005). *Border and transportation security: Possible new directions and policy options*. Washington, DC: Congressional Research Service, Library of Congress. Retrieved from http://www.fas.org/sgp/crs/homesec/RL32841.pdf

Royal Canadian Mounted Police. (2004). *Protection, detection & response* (Technical Security Branch Publication G1-025). Ottawa, ON, Canada: Author. Retrieved from http://www.rcmp-grc.gc.ca/ts-st/pubs/phys-sec/g1–025-eng.pdf

Schneier, B. (2000). *Secrets & lies: Digital security in a networked world*. New York, NY: Wiley.

Schneier, B. (2003). *Beyond fear: Thinking sensibly about security in an uncertain world*. New York, NY: Springer.

Talbot, J., & Jakeman, M. (2009). *Security risk management body of knowledge*. Hoboken, NJ: Wiley.

Trusted Information Sharing Network. (2008). *Defence in depth*. Australian Attorney General's Department Trusted Information Sharing Network. Canberra, ACT, Australia. Retrieved from http://www.tisn.gov.au/www/tisn/rwpattach.nsf/VAP/(99292794923AE8E7CBABC6FB71541EE1)~SIFT-Defence-in-Depth-CIO±-±15±Oct±2008.pdf/$file/SIFT-Defence-in-Depth-CIO±-±15±Oct±2008.pdf

U.S. Coast Guard. (2002). *Maritime Strategy for Homeland Security*. Washington, DC: Author. Retrieved from http://www.uscg.mil/history/articles/uscgmaritimestrategy2002.pdf

U.S. Department of Defense. (2003). *Protection Joint Functional Concept*. Washington, DC: Author. Retrieved from http://www.dtic.mil/futurejointwarfare/concepts/jroc_protection_jfc.doc

U.S. Department of Homeland Security. (2007). *The National Strategy for Homeland Security*. Washington, DC: Author. Retrieved from http://www.dhs.gov/xlibrary/assets/nat_strat_homelandsecurity_2007.pdf

U.S. Department of Homeland Security. (2009). *National Infrastructure Protection Plan: Partnering to enhance protection & resilience*. Washington, DC: Author. Retrieved from http://www.dhs.gov/xlibrary/assets/NIPP_Plan.pdf

Viega, J., & McGraw, G. (2001). *Building secure software: How to avoid security problems the right way*. Boston, MA: Addison-Wesley.

Willis, H., Morral, A. R., Kelly, T. K., & Medby, J. J. (2005). *Estimating terrorism risk*. Santa Monica, CA: RAND Corporation. Retrieved from http://rand.org/pubs/monographs/2005/RAND_MG388.pdf