

# A New Risk Management Paradigm For Assessments and Evaluations of Information Assurance Systems

George Jelen

## Abstract

*Traditional risk management methods developed and practiced by the Information Assurance (IA) community have typically begun with an identification of system threats or vulnerabilities, followed by a careful assessment of the consequences and the likelihood of each of these threats or vulnerabilities. Such an approach, although valid, tends to be unnecessarily labor intensive and often delays the ultimate necessary focus on the threats, vulnerabilities and countermeasures that are most important.*

*This paper sets forth a different approach, establishing a connection between information and mission, and then focusing primary attention on those information system objectives that are critical to that mission. The approach presented draws upon a general security risk management process advocated by the Joint Security Commission (JSC), which had been tasked by the Secretary of Defense and the Director of Central Intelligence to study the U.S. Government's security policies and procedures [1], but it strengthens the connection to mission and applies the process specifically to information systems. The resulting methodology can serve as a useful frame work for a revised approach to IA assessments and system evaluations—in effect, a new paradigm.*

## Introduction

In an era in which the battle for information is becoming a dominant objective of both warfare and commerce, systems that store and process information have become the battleground over which future competitions and conflicts will be waged. Both sides seek to know and to control the information available to their adversary or competitor. Increasingly, the side that "wins" this war over information is the side that prevails in whatever form of conflict or competition it is engaged. To succeed in this information war, an organization must have, not only an effective means of obtaining useful data about the adversary's intentions and strategies, but an equally effective means of protecting its own as well. As a result, the security of information systems has assumed much greater importance.

At the same time, with the collapse of the former Soviet Union, the U.S. has entered a period in which government expenditures for security are being challenged. The changed world situation, and the reduced security threat that it has brought with it, have caused many to question the continued need for any security protection. This reduced threat, coupled with a need to trim expenditures generally, have led to a cutback in resources dedicated to security. This cutback, in turn, is placing new demands upon security communities—including the Information Systems Security (IA) community—to do more with less.

As expenditures for security are reduced, it becomes ever more important that the reductions are made intelligently and that the remaining available funds are applied where they can do the most good. It is necessary to distinguish between what really needs protection and what does not. Some mechanism is required that allows the prioritization of security needs according to some orderly process—a process that is rational, convincing, repeatable, and reasonably simple.

Such a process should weigh the importance of the information asset, the motivation and the capability of the adversary, the ease with which that adversary could exploit or harm that asset, and the risk of leaving the asset unprotected versus the cost of protecting it. This is the role of the Information Systems Security (IA) assessment or system evaluation. The assessment or evaluation

must employ a reasonably rigorous security risk management paradigm that can be effectively applied to resource decisions.

Without such a rigorous application of risk management principles, any effort at assessing the security posture of an information system quickly reduces to little more than a search for matches between vulnerability and threat. Unless such a search is accompanied by an analysis of consequences, a prioritization of threats and vulnerabilities, and some kind of cost-benefit analysis, the application of countermeasures becomes an exercise in risk avoidance.

## **The Process**

Traditional risk management methods developed and practiced by the IA community have typically begun in one of three ways. Some have started by looking first for system vulnerabilities, others have begun with an examination of threats to the system, while still others have proceeded from formal statements of security requirements<sup>1</sup>. Although all of these approaches can produce valid results, they all run the risk of missing entirely the components of the system on which the using organization most depends, and thus deserve its primary attention.

For example, a fairly typical approach to a threat or vulnerability-focused assessment would be to first identify the system threats or vulnerabilities and then perform a careful assessment of the consequences and the likelihood of each of these threats or vulnerabilities. A straightforward description of this traditional process is presented in the *Information Systems Security Engineering Handbook* [6]:

*The inputs to the Security Risk Assessment process are the threats to the systems and the weaknesses within it. This information is used to determine whether the weaknesses are exploitable—thus making them system vulnerabilities. An assessment is then made of the likelihood and consequences of those vulnerabilities being attacked.*

Such an approach, although valid, tends to be unnecessarily labor intensive and often delays the ultimate necessary focus on those threats, vulnerabilities and countermeasures that are most significant. This often leads to the expenditure of resources for countermeasures that fail to address the most critical system shortcomings. Yet, because the threats or vulnerabilities revealed by the process have been dealt with, an unrealistic and false sense of security frequently results. In particular, many such past efforts have focused too intently on malicious modification or compromise of data while ignoring even more serious vulnerabilities that make it possible to deliberately sabotage the entire system through hardware or software subversion. A statement in the Drake and Morris paper [4] would appear to be apt: "This traditional model is assumed to be complete in its ability to model all of the countermeasures and represent **all** of the loss [emphasis added]." As the authors go on to point out, "We have not found this to be the case."

This paper sets forth a different approach, establishing a connection between information and mission, and then focusing primary attention on information requirements and those information system objectives that are critical to that mission. The approach presented maps easily into the process advocated by the Joint Security Commission (JSC). This commission was appointed by the Secretary of Defense and the Director of Central Intelligence in June of 1993, and tasked to study the U.S. Government's security policies and procedures and to develop a new approach to security that would assure adequate protection but be simpler, more uniform, and more cost effective.

---

<sup>1</sup> For example, Otwell and Aldridge [2] is a vulnerability based methodology; Jaworski [3] as well as Drake and Morris [4] present threat based approaches, and Stevens and Weiner [5] is an example of a requirements based process.

The Commission's final report, entitled *Redefining Security* [1], specifically advocates a risk management approach patterned after the methodology which had evolved over the years in the Operations Security (OPSEC) community. Although not intended expressly for information systems, this methodology can serve as a useful framework for a revised approach to IA assessments and system evaluations.

As outlined in the Commission's report, Security Risk Management applies a five-step process. The JSC process consists of asset valuation and judgment about consequences of loss; identification and characterization of the threats to specific assets; identification and characterization of the vulnerabilities to specific assets; identification of countermeasures, costs, and tradeoffs; and risk assessment. The five-step construct is not intended to represent a series of sequential steps but merely a logical formulation. Indeed, these five steps are interdependent and often overlap.

## **Asset Valuation**

As defined by the Joint Security Commission [1], asset valuation involves determining what is to be protected and determining its value. Both of these determinations are important, but at least for information systems, it is usually more helpful to deal with them in the opposite order: the mission criticality, i.e. the value, of a system and its information should dictate what most needs protecting. In an information system, the asset to be protected could be information that is resident on the system, information in transit through the system, or even the system itself. What is important here is to link value with how important the system or its information is to the fundamental mission or purpose of the organization. The more the organization depends upon a system or its information, the higher its value. It may not be necessary to quantify this value: a qualitative ranking is often sufficient. But if quantification is sought, Peebles [7] presents one approach to placing a value on information.

Although all of this sounds rather straightforward and simple, it is where many assessments go wrong. There is a strong bureaucratic tendency, particularly within the national security community, to equate need for protection with level of classification. This reflects both an unjustifiable bias toward confidentiality protection as well as an obsolete paradigm based primarily on adherence to security rules and regulations.

Data sensitivity, by itself, is a very poor indicator of the risk the data faces and generally is not even a good indicator of its value. Yet security policy and the formal security requirements based on the policy are more likely to reflect data sensitivity than mission criticality. Stevens and Weiner [5] attempt to resolve the tension between policy and mission criticality by addressing both within the same process step. Step 2 of their six-step process is called "Characterize your system, software, and data." They suggest doing this by focusing on two things: "(1) an assessment of the relative importance of the system, software, and data to their users and organization; and (2) an identification of what types of information you are processing (e.g. unclassified, sensitive but unclassified, or classified)." This association might be relevant and useful if one were concerned only with confidentiality or privacy. But more often than not, protection against denial of service and unauthorized modification of data are far more mission critical than protection against unauthorized disclosure, and the IA risk most threatening to the mission of the organization may very well not be covered by any extant rules or regulations.

As a way of assuring proper focus and as a first step toward asset valuation, an IA assessment or evaluation might well begin with obtaining answers to the following questions:

1. What are the principal missions of the organization?
2. What information systems are essential to the carrying out of these principal missions?
3. What data do these principal missions depend upon for their success?

4. What information, were it to fall into the hands of the adversary or competitor, would be sufficient to tip the scales in the adversary's favor?

The answers to these four questions should permit the generation of a prioritized list of undesired consequences, which can then be used to guide the rest of the assessment or evaluation.

By seeking answers to these four questions first, the limited resources available to perform IA assessments or system evaluations are focused on the most important issues. They are not expended chasing down vulnerabilities that will ultimately be found to have little or no consequence—at least relatively. The last three of these questions are minor restatements of the three major goals or objectives of information systems security, i.e. availability, integrity and confidentiality (although they are usually presented in the opposite order.)

The relative importance of each of these three components of IA depends largely on the role that the information system serves and on circumstances—most notably timing. For example, the consequence associated with a system's nonavailability is closely related to recovery time. A system that must be restored within an hour after disruption represents, and requires, a much more demanding set of policies and controls than does a similar system that need not be restored for two to three days. Likewise, the consequence associated with the loss of confidentiality can vary with time. For example, in a military engagement, advance knowledge of the battle plans of an attacking force could be gravely damaging if compromised well before the invasion. Once the invasion has begun, however, such information quickly loses its value. Similarly, in a business situation, early disclosure of a major product announcement may jeopardize competitive advantage, but disclosure just before the intended announcement may be insignificant. In both cases, although the information remains the same, the timing of its release significantly affects the consequence of the loss [8].

Sometimes, the systems most critical to an organization's mission are ones outside the organization's control. Most organizations, for example, are strongly dependent upon the civilian infrastructure. The Joint Security Commission [1] recognized this point:

*If, instead of attacking our military systems and data bases, an enemy attacked our unprotected civilian infrastructure, the economic and other results could be disastrous. Over 95 percent of Defense and Intelligence Community voice and data traffic uses the public phone system. The economic consequences alone of a successful attack on the phone system or the National Information Infrastructure would be significant.*

In spite of this strong dependency, these systems probably lie outside the boundary of our assessment or system evaluation. Nevertheless, the extent of this outside dependence must form part of the context in which any proposed corrective action is considered.

It is useful, when addressing the question of value, to view the question from the adversary's perspective as well as our own. As Smith [9] put it, "Asset attractiveness to the threat agent is different from asset value to the organization, reflecting the different value structure of the threat agent . . ." The significance of this point is that the value ascribed to an asset by the adversary is likely to determine the zeal with which that adversary pursues that asset. Because asset valuation needs to be adversary-centered, it can not be conducted without regard to analysis of the threat. Thus asset valuation, the first step in the process, and threat analysis, the second step, must be closely linked.

## **Threat Analysis**

The assessment process must account for two kinds of threat, i.e. inadvertent and deliberate. Inadvertent threats are those that derive from natural events, system malfunctions, or human error. Deliberate threats are those that are planned and perpetrated by individuals or

organizations bent on compromising, altering, impairing or destroying information assets. Too often, our traditional IA assessments or evaluations have ignored or played down inadvertent threats and have not even considered their impact or likelihood.

An analysis of inadvertent threats must first consider both the consequences of the event and its likelihood. If either consequence or likelihood is judged to be minimal or negligible, no further analysis of that particular threat is warranted. Only if both consequence and likelihood are judged to be at least moderate, would an examination of the system's vulnerability to such threats be in order.

Although the assessment process for deliberate threats could proceed in a similar manner, a more useful approach is usually through the development of the adversary's best and most likely strategy. For the purposes of this analysis, the assessment team assumes the point of view of the adversary, identifies its mission objectives, develops a set of strategies to achieve these objectives, and then selects the best or most likely of these strategies. Kerry [10] presents an excellent discussion of how the development of the adversary's strategy can be used to focus security assessment efforts.

Once the principal adversary's best and most likely strategy has been postulated, it becomes possible to focus protection resources in such a way as to defeat that strategy. The assumption is that the adversary or competitor is motivated to succeed in its strategy and could choose to do so either by preventing the friendly organization from achieving its fundamental objective or by besting it through superior information. Since the selection of the best or most likely strategy is greatly influenced by perceived vulnerabilities, the teasing out of specific threats must proceed in parallel with an analysis of vulnerabilities.

The Air Force Information Warfare Center at Kelly AFB further categorizes deliberate threats as either systematic or nonsystematic. In their formulation, non-systematic threats are those originating from individuals acting alone or from nonprofessional, minimally-resourced groups or organizations. Systematic threats are those from foreign governments or well-heeled organizations such as drug cartels or well-supported terrorist organizations. Most organizations, be they government or commercial, are likely to be subject to both systematic and non-systematic threats.

## **Vulnerability Assessment**

The third step in the assessment process is a search for vulnerabilities. This search should be guided by the prioritized list of undesired consequences mentioned earlier. If the resulting consequence is not deemed to be serious, then any vulnerability that produces it is not serious either.

In assessing vulnerabilities, it is again useful to assume the point of view of the adversary. Vulnerability assessments are intended to identify weaknesses in an asset that the adversary can exploit. We take another look at the adversary's strategy generated in Steps 1 and 2. With an understanding of what the adversary hopes to accomplish, this step looks at how the strategy might be carried out. Are there some system weaknesses that could be exploited in such a way as to advance the adversary's strategy?

When evaluators and assessors adopt the adversary's point of view, they inevitably become much more creative in their search for vulnerabilities. Lacking this adversarial perspective, one is likely to conduct the search for vulnerabilities through an assessment of each of the system's built-in safeguards, and perhaps even to allow oneself to be impressed by the strengths of some of these features. But, when dealing with security systems, it is important to keep in mind that the strongest features of the system are the least relevant to a vulnerability assessment. What one needs to look for and concentrate on are the weaknesses of the system. Assuming the adversary's point of view helps to assure this perspective.

## Countermeasures, Costs and Tradeoffs

The product of the previous step is a set of vulnerabilities, each of which could result in a serious consequence. Because of the initial focus on mission criticality, those not capable of producing a serious consequence were never considered.

Typically, in any given information system, there are many ways in which an information system could be impaired or its data modified or compromised. But not all of these vulnerabilities are worth correcting. For example, if all of these various ways were to be ranked as to their ease or likelihood, it makes little sense to spend much money to correct the fifth vulnerability on the list if there is nothing that can be done about the second.

For every identified vulnerability, there is usually a fairly obvious countermeasure—a way of countering that vulnerability. But some of these countermeasures may be expensive—either in monetary or in operational terms. Before blindly embarking upon a program to implement a set of countermeasures, it is important that their costs be identified. Sometimes, too, different countermeasures are of comparable effectiveness in countering a given vulnerability. In such cases, it obviously makes sense to implement the cheapest and easiest. Thus, relative costs and impacts of alternative countermeasures should be evaluated before implementing any of them.

The assessment process also needs to consider combinations of protective measures. The report of the Joint Security Commission [1] talks about "an approach to . . . information systems security [that] is seen as part of a balanced mix that also includes personnel security, physical security and other security procedures." Specifically, the report states:

*Security must come from an integrated system that recognizes the interdependence of the individual security disciplines and establishes a logical nexus between the sensitivity of information and the personnel, physical, information, and technical security countermeasures applied in protecting the information.*

## Risk Assessment

The previous step should have yielded a ranking of potential countermeasures according to their expected payoff. However, just because a given countermeasure is found to be the most cost effective does not necessarily mean that it deserves to be implemented. The benefit derived must at least equal the cost. Even the least expensive countermeasure may not be worth the cost to implement. Sometimes too, the elimination of one vulnerability introduces another. It is at this point in the process that risk assessment comes into play. In this step, the cost of implementing a given countermeasure is compared with the cost or loss associated with not implementing it. The cost of not implementing a countermeasure and therefore of leaving a vulnerability uncorrected is generally called "risk." Risk is a function of both threat and vulnerability: if either threat or vulnerability exists without the other, there is no risk. A proper assessment of risk includes a cost/benefit analysis of proposed countermeasures that independently assesses the gain versus the loss resulting from the elimination of each vulnerability. Of course, it is impossible to eliminate risk entirely; one can only mitigate or reduce it.

Finally, the analysis needs to recognize when enough is enough. As the JSC [1] observed, "In many cases, there is a point beyond which adding countermeasures will raise costs without appreciably enhancing the protection afforded." The evaluating or assessing organization does damage to its credibility when its recommended countermeasures push beyond the point of diminishing return.

## **Importance of the Five Steps**

Each of the five steps is important to the integrity and efficacy of the overall process. Although each of them has value in and of itself, it is only when all five are employed together that the full synergistic value of the risk management process accrues.

Asset valuation provides focus. If an organization does not take the time and make the effort to identify the systems and information most critical to its mission, with the intent to focus scarce protection resources on such systems and information, an organization can easily find itself trying to protect all systems and all information equally. This makes no sense and no organization can afford to do so even if it did.

Threat analysis assures realism. Just as it is not possible to protect all information, it is also not possible to protect any information against all threats. To be worth protecting against, a threat has to be realistic. Competent threat analysis assures this.

Vulnerability analysis lends objectivity. The critical element of a vulnerability analysis is placing ourselves in the position of the adversary as we probe our information systems in depth, searching for ways that our adversary could exploit or impair them. Looking at our own systems through the eyes of our opponent is what brings objectivity as well as imagination to the effort.

The final two steps, cost and tradeoff analysis of countermeasures, and risk assessment, guarantee rationality. Fundamental to the process is the tenet that not all vulnerabilities are worth correcting. Deciding which ones are and which ones are not is the task of the last two steps in the methodology. Thus these final two steps are key in ensuring that the overall process succeeds as a means of managing risk. Absent a fairly rigorous assessment of the risk, the rationale necessary for intelligent countermeasure decisions is missing.

Taken together, these five steps do indeed represent a logical and balanced approach to contending with risk. The point, however, is that one has to complete all five of them. Because the asset valuation step is what lends focus, if that step is omitted, the analysis is unfocused. Similarly, if the threat analysis step is dropped, the process lacks realism. Absent an adversarial analysis of vulnerabilities, it would not be objective. And, if we skip the countermeasures tradeoff and risk assessment, the effort lacks rationality. The real power of the process results from the combined contribution of all five steps. Leaving any one of them out robs the process of that power. The report of the Joint Security Commission [1] makes the same point. Commenting on their process, their report notes: "When any of the steps are left out, the result can either be inadequate protection or unnecessary and overly expensive protection." Particularly in an era of scarce resources, either of these outcomes is unacceptable.

## **A New Paradigm**

For some security people, the risk management process may require a whole new way of thinking in other words, a paradigm shift. The old paradigm has tended to look at the problem from the inside out; the new paradigm looks at the problem from the outside in. An inside-out look begins with an internal examination of the system by identifying requirements or searching for vulnerabilities; an outside-in approach begins with an examination of the operational environment within which the system must perform. As this paper has attempted to show, there are at least two advantages in viewing the problem from the outside. The first is that it allows one to step back from the problem and ask what is it that the customer or owner of the system is most concerned about. What events, what circumstances, which information losses would impact most severely on the ability of the customer organization to perform its mission? These questions are rarely asked with the result that whatever analysis takes place is too quickly focused on some esoteric, technical problem that even if corrected will matter little. The second is that it permits one to look at the problem from the point of view of the adversary or competitor. What information is most critical to the success of the adversary's strategy? It is this information that the

adversary or competitor is most likely to go after and it is this information that most needs protecting. And finally, it conserves resources by avoiding spending time or effort exploring vulnerabilities of little consequence.

In addition to requiring a different perspective, the new paradigm is directed toward a different objective. Whereas the old paradigm attempted to enforce a set of rules and regulations, and protect only what is legally required (which partially accounts for the fixation on confidentiality protection), the new paradigm leaves much room for judgment. It attempts to teach people how to discern and to make thoughtful choices. As the JSC [1] put it:

*Conceptually, [security] should be the way we think rather than a manual of rules. Security then becomes a more positive undertaking that values the spirit over the letter of the law, problem prevention over problem resolution, and individual responsibility over external oversight.*

Clearly, risk management demands that we weigh costs and benefits. It requires us to make reasoned choices. But it requires more than that. First and foremost, it requires that we protect what is most important and not waste scarce security resources on protecting that which, at least comparatively, is either of little consequence or is already lost.

Too much money has been and continues to be spent trying to protect information that is already in the public domain. I am not suggesting here that we give it away or confirm it just because it appears in open literature; but neither does it make much sense to expend increasingly hard-to-get security resources in a futile effort to protect it. Here, the discipline of Operations Security (OPSEC) could be helpful, since part of the OPSEC methodology is to assess what is in the public domain.

Given the increased openness of our society, there are many who would question whether or not efforts to protect any information are even worth the struggle. Already people question employing intelligence resources to collect technology-related information. They argue that, in the West, since governments are so open, presses so free, and information so readily available, using scarce intelligence resources to gather what is easily obtained from the media, journals or open dialogue seems unnecessary and wasteful. But I would observe that, if it is pointless to use scarce intelligence resources to collect it, it is at least equally pointless to dedicate scarce security resources to protect it. The judicious application of the risk management approach espoused herein can guard against this.

## **The Insider Threat**

Frequently, the greatest threat to an information system comes from an insider—someone who has been granted legitimate access to the information system. This is particularly insidious because the threat source may be acting completely within his or her authorizations. Of course, the fact that a given individual has the administrative authority to commit mischief renders the consequences of such mischief no less of a problem. But it undoubtedly makes the problem much harder to solve.

An important consideration in evaluating the extent of the insider threat is a determination of the probability that a "bad guy" is among the set of authorized users. Historical data relevant to this determination can often be obtained from the organization's security or counterintelligence department.

## **A Defensive Activity**

Information Systems Security is a defensive activity,<sup>2</sup> aimed at thwarting the offensive efforts of someone else. Offensive disciplines, i.e. forms of information collecting or disrupting, always enjoy an advantage over defensive disciplines like IA. The offense only has to succeed once; the defense has to succeed each and every time. The offense has only to find one offensive strategy that works; the defense has to defend against all offensive strategies. Furthermore, the offense usually learns immediately when it succeeds; the defense never knows whether it is succeeding or not. The defense occasionally finds out when it has failed but typically this is many years after the fact.

As with any other defensive activity, IA's major goal is to foil the adversary's offensive strategy. Obviously, in order to do this, one has to give considerable thought to what that strategy is likely to be. Since defensive measures must be "prepositioned," i.e. already be in place at the time the offense carries out its strategy, the offense always has the advantage of the last move. In the series of moves and countermoves that characterize the continual battle for technological dominance between offense and defense, the offense always retains this large advantage. It is the offensive side that controls the timing of the actual engagement. The offense, therefore, is always able, to optimize both the timing and the attack strategy against a generally known defense. Although there may exist some adequate defense for every specific attack strategy, there will be no time, in actual conflict, for the defensive side to field the appropriate countermeasure once the engagement has begun. [11]

## **Assessment Results**

There are many ways in which the results of the assessment or evaluation can be communicated in a final report. e.g. risk calculations, figures of merit, etc. All have their utility and their proponents. Regardless of the means used to convey them, the results must be supportable, credible and balanced. The final report should include a recommended set of countermeasures aimed at managing risk—not avoiding it, which is impossible anyway.

In whatever form employed to communicate the results, assessments and evaluations can be very useful. They can serve as the first step toward the development of a defensive strategy. And in this day and age, no organization should be heading into its information war without one.

## **Acknowledgements**

By sharing their insights and refining mine, many persons, have contributed to this paper and I acknowledge my indebtedness to all of them. I thank, in particular, all those in the IA community who taught me that it is far easier to penetrate a computer than to secure one, and the many in the OPSEC community from whom I learned the importance of a mission focus. Finally, I wish to thank the reviewers of my original paper who offered a number of helpful suggestions, most of which I tried to apply. I believe that my paper is much improved as a result.

---

<sup>2</sup> Intelligence gathering, and specifically Signals Intelligence or SIGINT, is the counterpart offensive activity.

## **References:**

1. U.S., Joint Security Commission, *Redefining Security: A Report to the Secretary of Defense and the Director of Central Intelligence*, 28 February 1994.
2. Ken Otwell and Bruce Aldridge, "The Role of Vulnerability in Risk Management," *Proceedings of the Fifth Annual Computer Security Applications Conference*, December 1989, pp. 32-39.
3. Lisa M. Jaworski, "Tandem Threat Scenarios: A Risk Assessment Approach," *Proceeding of the 16th National Computer Security Conference*, Baltimore, MD, September 1993, pp. 155-164.
4. David L. Drake and Katherine L. Morse, "The Security Specific Eight Stage Risk Assessment Methodology," *Proceeding of the 17th National Computer Security Conference*, Baltimore, MD, October 1994, pp. 441-450.
5. Jennie A- Stevens and Richard E. Weiner, "A Structured Approach to Risk Assessment: An Innovative Concept," *Proceedings of the 12th National Computer Security Conference*, Baltimore, MD, October 1989, pp. 472-482.
6. U.S., National Security Agency/Central Security Service, *Information Systems Engineering Handbook, Release 1.0*, February 28, 1994.
7. Donald R. Peeples, "Value of Information," *International Security Systems Symposium and Exposition Conference Proceedings*, Washington, DC, November 15-17, 1993, pp. 1-4.
8. National Research Council, *Computers at Risk: Safe Computing in the Information Age*, Washington. DC: National Academy Press, 1991.
9. Suzanne T. Smith, "LAVA's Dynamic Threat Analysis," *Proceedings of the 12th National Computer Security Conference*, Baltimore, MD, October 1989, pp. 483-494.
10. George T. Kerry, "The OPSEC Survey: Start to Finish," *Proceedings of the Third National Operations Security Conference*, Las Vegas, NV, May 1992, published by the Interagency OPSEC Support Staff.
11. George F. Jelen, "Space System Vulnerabilities and Countermeasures," in William J. Durch, ed., *National Interests and the Military Use Of Space* (Cambridge, MA: Ballinger, 1984), pp. 89-112.