CONF-871108--4

UCRL 97741
PREPRINT

UCRL--97741

DE88 008677

Safeguards Effectiveness Evaluations In Safeguards Planning

Rokaya A. Al-Ayat

This paper was prepared for presentation at the American Nuclear Society Conference November 30-December 4, 1987 San Diego, California

December 3, 1987

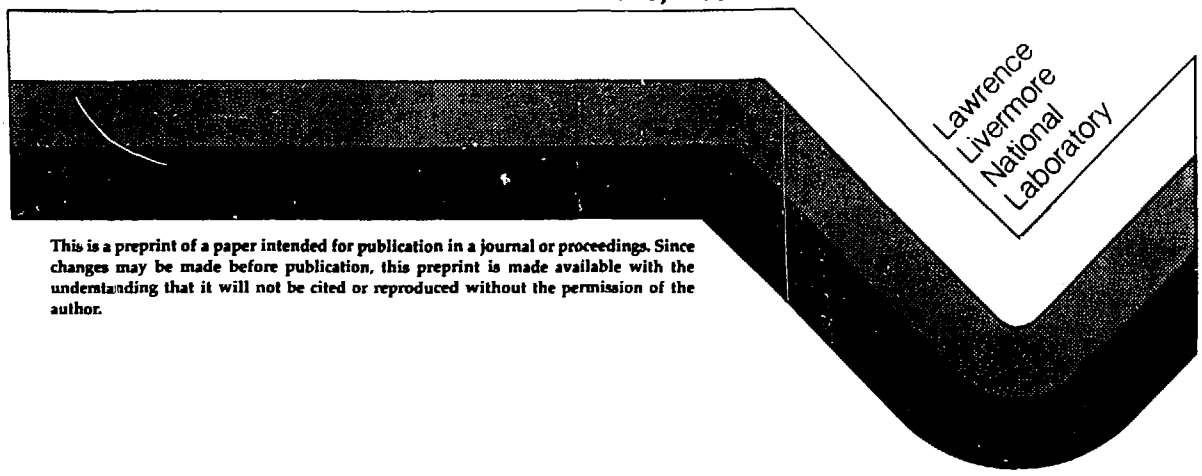

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

SAFEGUARDS EFFECTIVENESS EVALUATIONS IN SAFEGUARDS PLANNING*

ROKAYA A. AL-AYAT, Lawrence Livermore National Laboratory
7000 East Avenue
Livermore, California 94550
415-422-8467

ABSTRACT

This paper describes analytic tools we developed to quantify the effectiveness of safeguards against theft of special nuclear material by insiders. These tools help identify vulnerabilities in existing safeguards, suggest potential improvements, and help assess the benefits of these upgrades prior to implementation. Alone, these tools are not sufficient for safeguards planning, since the cost of implementing all suggested upgrades almost always exceeds the available resources. This paper describes another tool we developed to allow comparison of benefits of various upgrades to identify those upgrade packages that achieve the greatest improvement in protection for a given cost and to provide a priority ranking among cost-effective packages, thereby helping decision-makers select the upgrades to implement and highlight the amount of residual risk.

INTRODUCTION

This paper discusses safeguards effectiveness evaluation techniques and the role they play in safeguards planning. It is a challenging problem to plan an effective and balanced safeguards system that adequately protects facilities handling special nuclear material from malevolent acts. First, safeguards planners must consider a variety of hypothetical threats as specified in generic threat guidance. These threats include terrorists, criminals, and insider adversaries. The latter threat poses particular difficulties because insider adversaries have access to protected assets and knowledge of operations and safeguards. Moreover, protection against insider attempts almost always constrains production and can, therefore, be costly. Second, many of the safeguards options available have varying levels of success against the various threat types. These options include physical security measures, material control, material accountability, and human reliability programs. Third, safeguards that may be effective against one threat type may be detrimental to other threat types. Because of the complexity of the problem, analytical tools are needed to evaluate safeguards effectiveness, to assist in setting upgrade priorities, and to allocate limited safeguards resources.

We have developed several systematic and quantitative methods for evaluating the effectiveness of existing safeguards and proposed modifications. These methods vary from very detailed -- designed to evaluate safeguards effectiveness when confronted with every conceivable adversary action -- such as the Structured Assessment Approach (SAA)[1] -- to simple, straightforward methods that identify weaknesses and pinpoint areas where a more detailed analysis may be warranted. At this time, the most widely used tool for insider threats is the Safeguards Evaluation Tool (ET).[2] A similar tool was developed by Sandia National Laboratories for identifying vulnerabilities against terrorist attacks. This tool is known as SAVI (Systematic Analysis of Vulnerability to Intrusion).[3]

Safeguards evaluation methods, such as ET and Sandia's SAVI, generally identify vulnerabilities in existing safeguards and suggest potential improvements against insiders and outsiders. Alone, they are not adequate for safeguards planning because the cost of implementing all suggested upgrades generally exceeds the available resources. This paper describes a tool we developed to integrate the results of an insider and outsider threat evaluation, to allow comparison of benefits of various upgrades, and to identify those upgrade packages that achieve the greatest improvements in protection for a given cost. The tool, MI$ER[4] also provides a priority ranking among the cost-effective packages, thus allowing a decision-maker to decide which upgrades to implement and the amount of residual risk to accept.

Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

We describe the philosophy and approach of SAA, ET, and MI$ER. Then we conclude with a very brief discussion of our current effort to develop an integrated vulnerability assessment tool that can handle insiders, outsiders, and insiders and outsiders in collusion.

EVALUATING THE EFFECTIVENESS OF SAFEGUARDS

Protecting against the threat posed by insiders is a particularly difficult problem because such adversaries have routine access to assets, as well as knowledge of the operations and safeguards that are in place. Furthermore, protection against insider theft almost always constrains normal operations. Although no terrorist attacks have occurred in DOE facilities or those licensed by the Nuclear Regulatory Commission (NRC), insider events have ranged from rather inconsequential theft of property to infrequent but more damaging attempts at espionage.

Potential insider adversaries include anyone with access to assets that are safeguarded. For example, for a facility that manufactures research-reactor fuel containing special nuclear material (SNM), potential adversaries include process operators and supervisors, nuclear material custodians and accountants, janitors, security inspectors, computer programmers, and others who may be acting either alone or in collusion with other adversaries.

Protection against such threats requires an integrated system of safeguards consisting of human reliability programs, physical protection, material control, and material accountability. The first of these measures, human reliability programs (HRPs), includes security clearances, security awareness activities, and psychological screening programs, all of which are designed to deter or reduce the likelihood of malevolent acts. Because HRP measures have limitations -- and an otherwise reliable insider may be coerced -- total reliance cannot be placed on them. Facilities therefore use other measures to detect potential theft or diversion attempts not deterred by HRPs.

The second protective measure is physical protection, such as containment and controls that limit access to protected assets. However, insiders generally have authority to override these measures. The third measure, which plays an important role in protection against SNM theft, is material control and accountability (MC&A). MC&A includes SNM monitors on doors to areas where material is stored or processed, as well as periodic physical inventories. However, these systems may be subject to tampering or inaccuracies. MC&A systems rely heavily on administrative procedures that can be circumvented if they are not designed properly. Moreover, records and forms are subject to inadvertent errors and falsification, thereby decreasing confidence in them. Evaluating the effectiveness of these protection measures against insider threat is a complex task.

We have developed a series of systematic and quantitative methods to assist decision-makers in evaluating the effectiveness of their safeguards systems against theft or diversion of SNM by insiders. These methods vary widely in complexity. The most detailed ones, such as the Structures Assessment Approach (SAA), are designed to evaluate safeguards effectiveness with respect to every conceivable adversarial action and plant operating condition. At the other extreme are simple, straightforward methods that identify weaknesses and pinpoint areas needing more detailed analysis. At present, the most widely used tool in the latter category is the Safeguards Evaluation Tool (ET), which is currently used for insider threats. Although ET is based on more complex models developed at LLNL, we have refined it to a straightforward tool for use by safeguards and security planners at their facilities.

Structured Assessment Approach (SAA)

We have developed the SAA to help assess the vulnerability of safeguards systems to insiders. For physical security systems, the SAA identifies possible paths that are not safeguarded under various operating conditions, including day or night shifts and emergency conditions. SAA also identifies the insiders who could defeat the system via direct access, collusion, or indirect tampering. For MC&A systems, this method also identifies personnel who could block the detection of lost or diverted material by falsifying data or tampering with equipment.

We first used the SAA in late 1978 to assess a nuclear fuel facility; we have subsequently applied it at several DOE and NRC facilities. Although the SAA was originally designed to run on a mainframe computer, we recently converted it to a user-friendly program that runs on a personal computer. We have added many features to simplify and facilitate its use in conducting vulnerability analyses. For example, the SAA input is a text-like data file that is easily read and can provide documentation of both the facility safeguards and the assumptions used during analysis. However, because of the level of detail required for an SAA analysis, this tool continues to be used primarily by analysts at LLNL.

Safeguards Evaluation Tool (ET)

Although ET is a relatively simple tool to use, it captures the essence of most complex models. ET allows the user to generate quantitative measures of performance that can be used to judge the effectiveness of systems against

various threats, to pinpoint vulnerabilities, and, most important, to compare the performance of alternative system upgrades. We specifically designed ET for safeguards and security planners who are interested in evaluating their own facilities. ET can also be used by appraisal teams as they visit various sites. The method has been applied successfully at many DOE facilities.

We have developed an Evaluation Workbook to guide an appraisal team through the steps involved in data collection and evaluation. The Workbook also provides a means for documenting the characteristics of a given safeguards system and the assumptions that are made during an evaluation. The descriptions of threats include a list of potential adversaries, their authority, their access to SNM, and the number of each type of potential adversaries in the facility. Safeguards components are also recorded, especially those designed to limit access to SNM or unauthorized movement of SNM. Because many of the safeguards designed to protect assets against insiders are procedural (for example, the two-person rule), it is important to document these procedures and the details of their day-to-day implementation at the facility.

ET allows an evaluation team to quantify the effectiveness of an integrated system of physical security and MC&A against the insider threat. ET uses the probabilities of timely and late detection as measures of safeguards effectiveness against theft or diversion of SNM. Detection is timely if a theft is discovered before material is removed from the site; detection is late if the discovery is made after the material has been removed.

To simplify the evaluation of timely detection, a theft attempt is divided into three stages:

- SNM acquisition, which includes gaining possession of SNM inside a material access area and concealing it for later removal.
- Removal of material from a material access area, which includes transporting concealed SNM to a location outside that area and hiding it in a protected area for later removal.
- Removal of material from a protected area to an off-site location.

Because diversion is the removal of material from its authorized location, it can be considered equivalent to the acquisition stage. A theft is successful when an adversary completes all three stages.

Once the evaluation team collects and documents threat and safeguards information, they use this data to identify the various paths and strategies each adversary could use at each stage of an attempted theft. In developing the strategies, we assume that an adversary would use stealth and deceit. Adversaries would consider many factors, such as timing of the attempt and exploitation of special knowledge, access, or authority. Adversaries would also choose those conditions that minimize or delay the chances of detection, such as falsification of records or substitution of material.

Next, the user assesses probabilities of timely detection for each adversary using each possible strategy for each of the three stages of theft attempt. These probabilities can be assessed by a group of knowledgable individuals. They can also be based on experimental data. These probabilities are used as input to ET. Next, the user runs ET to compute the probability of timely detection with the strategies and safeguards that are in place. We assume that an adversary would choose the strategy with the lowest chance of detection. ET computes the overall probability of timely detection for each type of adversary and then generates both tables and graphs, such as the graph shown in Figure 1. These results make it easy to compare the effectiveness of existing safeguards against various adversaries and to identify those with low probabilities of detection. This reveals areas where the safeguards system may be weak and needs improvement.

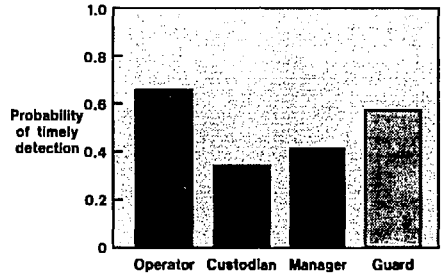

Fig. 1 The ET computer program computes and displays the overall probability of timely detection for each type of insider adversary.

Late detection, which by definition occurs after the fact, is important because it provides an indication that SNM may be missing. The sooner a loss is detected, the greater the possibility of resolving the cause, thereby preventing an incorrect response or assuring that no theft occurred. In case of an actual theft, quick detection and resolution will

hasten recovery of the SNM and mitigate adverse consequences.

Evaluation of late detection using ET is discussed in more detail in a companion paper in this issue. The reader is referred to the paper entitled, "Evaluating Late Detection Capability Against Diverse Insider Adversaries".[5]

Once an evaluation is complete, two additional steps are taken. The first step is to test the sensitivity of results to alternative assumptions and probability assessments. Many inputs for timely and late detection rely heavily on subjective judgments of the user, and ET provides a simple way to determine the effect of changing assumptions on the results. This sensitivity analysis highlights those inputs requiring investigation and focuses debate on the most important assumptions.

The second step involves using ET to identify those adversaries and stages for which the system is weak and to suggest areas where corrective actions or upgrades are required. ET can assess the change in safeguards effectiveness prior to implementation of upgrades so that the benefits of various configurations can be highlighted.

SETTING PRIORITIES FOR SAFEGUARDS UPGRADES

So far, the discussion has focused on the methods we developed at LLNL for evaluating safeguards effectiveness against the insider threat. A similar tool, developed at Sandia National Laboratories, Albuquerque, can identify vulnerability to terrorist attacks. This tool is known as the Systematic Analysis of Vulnerability to Intrusion (SAVI).

Both ET and SAVI provide systematic approaches for: identifying vulnerabilities, suggesting corrective actions or upgrades, and determining the effectiveness of proposed upgrades. However, neither ET nor SAVI helps decision-makers balance the effectiveness of proposed upgrades with their cost so that priorities can be set. We developed MI$ER (Method for Integrating SAVI and ET Results) to do exactly that.

THE MI$ER PROGRAM

MI$ER integrates the results of both insider and outsider evaluations, compares the benefits resulting from various upgrades, and identifies those upgrades that achieve the greatest improvement in protection for a given cost. MI$ER also ranks various packages of cost-effective upgrades, thereby helping decision-makers select the upgrades to implement and highlighting the amount of residual risk. Figure 2 is a typical MI$ER plot showing how the effectiveness of safeguards can change as different upgrades are implemented. This kind of analysis can also identify those upgrades that do not increase system effectiveness. Such a result can occur when an upgrade package designed to eliminate vulnerability against one threat introduces a vulnerability to another threat.

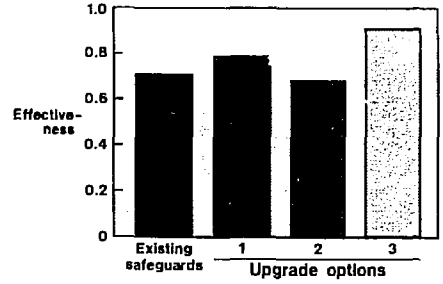

Fig. 2 The MI$ER program can be used to compare the change in safeguards effectiveness as different packages are implemented. Such comparisons can help decision-makers determine which upgrades should be implemented and which may actually introduce a new vulnerability to another threat.

Figure 3 is a MI$ER cost-benefit graph highlighting those upgrade packages that provide the most benefit per dollar and those with marginal return. The curve connecting the packages that offer the most benefit for a given cost is called the cost-effective frontier. Upgrades below the cost-effective frontier are generally inferior because they represent lower effectiveness per dollar spent. Such graphs provide decision-makers with a rational basis for choosing upgrades and for deciding what level of residual risk to accept.

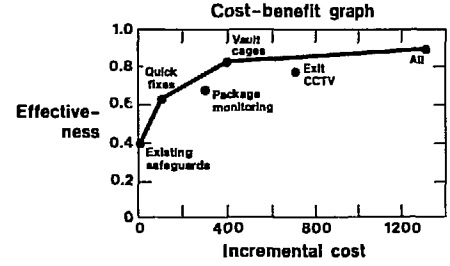

Fig. 3 A MI$ER cost-benefit graph is used to highlight those upgrade packages that provide the most benefit per dollar spent. Here, the solid line connecting high-priority packages is called the cost-effective frontier. Upgrades falling below this line, such as closed circuit TV here, are generally inferior because they represent lower effectiveness per dollar spent.

Although MI$ER provides the basic framework for setting priorities, significant additional research is still needed. MI$ER currently focuses on a single target, whereas we often need to allocate funds to multiple targets and facilities. In addition, the current model was designed for SNM theft, and this approach needs to be extended to other areas, such as radiological and industrial sabotage. These extensions will make MI$ER appropriate for a wide range of users including: personnel at DOE headquarters who must allocate limited funds to DOE weapon facilities; managers who are responsible for determining the relative benefits of new safeguards; and facility managers who must obtain the best possible protection with their limited safeguards resources.

CURRENT DEVELOPMENT

We're working with Sandia National Laboratories to develop an integrated vulnerability assessment program for addressing insider and outsider threats. The integrated program will enable the user to evaluate safeguards effectiveness against a variety of threats including: terrorists, criminals, demonstrators, single and colluding non-violent insiders, and -- to a limited extent -- insiders in collusion with outsiders. The program is integrated in the sense that the user will be required to describe the facility safeguards only once. The user will then be able to use this description for either the insider or outsider module. Results of these evaluations can then be used to assess effectiveness against collusion among insiders and outsiders.

There will be many differences between the new insider model and ET: the new evaluation model will contain data bases of adversary attributes, strategies, and baseline probabilities of detection. In addition, probabilities of detection will depend upon the safeguards present and their implementation, as well as the access and authority of each insider adversary. Also, the new outsider model will use a faster algorithm than SAVI and will model a larger spectrum of outsider adversaries. A model to calculate the probability of neutralization will be included.

The integrated package is due for release in 1988, and it is expected that, after the release, a training program in the use of the package will be developed and offered at the DOE Central Training Academy in Albuquerque, New Mexico.

CONCLUSIONS

Our safeguards evaluation tools -- such as SAA and ET -- and priority-setting tool MI$ER have benefited the safeguards community by enhancing the level of protection against potential insider and outsider threats, making cost-effective use of resources, and ensuring the consistency and comprehensiveness of safeguards and security policy. We continue to refine our evaluation tools and to transfer them to other facilities and institutions.

We are working with analysts from Sandia National Laboratories to develop an integrated package that can handle threats from insiders, outsiders, and insiders and outsiders in collusion. In fiscal year 1988, we expect to release a preliminary version of this package, which will include more powerful and updated versions of ET and SAVI.

REFERENCES

1. C. J. PATENAUDE, A. SICHERMAN, AND I. J. SACKS, "The Structured Assessment Approach: A Microcomputer-Based Insider-Vulnerability Analysis Tool," presented to the Institute of Nuclear Materials Management 27th Annual Meeting, New Orleans; June 1986, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-4265 (1986).

2. A. E. WINBLAD, "The SAVI Vulnerability Assessment Model," Institute of Nuclear Materials Management, 28th Annual Meeting, Newport Beach, California, July, 1987.

3. R. A. AL-AYAT et al., "Safeguards Evaluation Method -- Insider Threat," Lawrence Livermore National Laboratory, Livermore, CA, UCID-20145 (1986).

4. R. A. AL-AYAT, B. R. JUDD, and A. SICHERMAN, "Setting Priorities for Safeguards Upgrades," Institute of Nuclear Materials Management, 28th Annual Meeting, Newport Beach, California, July, 1987.

5. A. SICHERMAN, Evaluating Late Detection Capability Against Diverse Insiders, Lawrence Livermore National Laboratory, Livermore, CA, UCRL-97742 (1987).