SAND2013-5095C

Risk-Informed Management of Enterprise Security: *Methodology and Applications for Nuclear Facilities*[*]

Felicia A. Durán[†], G.D. Wyss, S.E. Jordan, and B.B. Cipiti

Security Systems Analysis
Sandia National Laboratories
Albuquerque, New Mexico
United States

Abstract. Decision makers wish to use risk analysis to prioritize security investments. However, understanding security risk requires estimating the likelihood of attack, which is extremely uncertain and depends on unquantifiable psychological factors like dissuasion and deterrence. In addition, the most common performance metric for physical security systems, "probability of effectiveness at the design basis threat" [*P(E)*], performs poorly in cost-benefit analysis. This makes it difficult to prioritize investment options on the basis of *P(E),* especially across multiple targets or facilities. To overcome these obstacles, work at Sandia National Laboratories has developed a risk-informed security analysis method. This methodology, Risk-Informed Management of Enterprise Security (RIMES), characterizes targets by how difficult it would be for adversaries to exploit each target's vulnerabilities to induce consequences. Adversaries generally have success criteria (e.g., adequate or desired consequences and thresholds for likelihood of success), and choose among alternative strategies that meet these criteria while considering their degree of difficulty in achieving their "successful" outcome. RIMES has been applied to evaluate the theft and sabotage risks for two types of nuclear fuel cycle facilities – used nuclear fuel (UNF) storage and small modular reactors (SMRs).

1. Introduction

Over the last several decades, security analysts have developed robust design processes and evaluation tools [1] to ensure that a fielded security system will provide effective protection against the adversaries for which it was designed. Sophisticated tools help analysts evaluate possible attack paths [2] and predict the effectiveness of security responders [3] so that decision makers can have reasonable assurance that the design-basis adversaries would likely be defeated should they attempt an attack. These systematic analyses ensure the security of many important assets and system. [4]. The chief metric for this assurance is the "probability of effectiveness for the design basis threat (DBT)" (P_E for DBT, or $P_{E|DBT}$), which represents the probability that a design basis adversary will fail to achieve the attack objectives even if the most advantageous attack scenario were attempted.

For many years, safety investment decisions have been made using risk-informed cost-benefit analysis in which the benefit metric is heavily based on a quantitative estimate of risk reduction. Many seek to perform similar analyses to prioritize security investments. However, understanding security risk requires estimating the likelihood of attack, which is extremely uncertain and depends on unquantifiable psychological factors like dissuasion and deterrence. In addition, $P_{E|DBT}$ performs poorly in cost-benefit analysis. It is extremely sensitive to small changes in adversary characteristics when the threat is near a system's breaking point, but very insensitive to those changes under other conditions. This makes it difficult to prioritize investment options on the basis of *P(E),* especially across multiple targets or facilities.

To overcome these obstacles, work at Sandia National Laboratories has been done to develop a risk-informed security analysis method. This methodology, Risk-Informed Management of Enterprise Security (RIMES), characterizes targets by how difficult it would be for adversaries to exploit each target's vulnerabilities to induce consequences. Adversaries generally have success criteria (e.g., adequate or desired consequences and thresholds for likelihood of success), and choose among

[*] Sandia National Laboratories is a multi program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000. Approved for Unclassified/Unlimited Release.

[†] Sandia National Laboratories, P.O. Box 5800 MS-0757, Albuquerque NM 87185-0757, United States

alternative strategies that meet these criteria while considering their degree of difficulty in achieving their "successful" outcome. Investments reduce security risk as they reduce the severity of consequences available and/or increase the difficulty for an adversary to successfully accomplish their most advantageous attack. To apply this insight requires development of a robust metric to characterize targets in terms of an adversary's degree of difficulty to prepare for and execute successful attacks. This requires one to compare and aggregate the relative difficulty for disparate adversaries to successfully acquire the requisite resources and employ them against specific targets. An objective risk-informed method can be developed from this metric, and applied to security investment prioritization using traditional optimization algorithms. A focus on the level of difficulty of a particular attack as opposed to the probability of attack will enable decision makers to balance competing security interests (e.g., multiple facilities) and provide objective and unbiased justification for investment decisions, resulting in more robust and cost-effective security systems. This shift allows for designers to manage risk better by balancing increased security against those threats that require lower difficulty for an adversary to produce higher consequences.

Most recently, RIMES has been applied to evaluate the theft and sabotage risks for two types of nuclear fuel cycle facilities – used nuclear fuel (UNF) storage and small modular reactors (SMRs). This paper discusses the development of the RIMES method and summarizes it application for uNF storage and SMR security.

2.0 Probabilistic Risk Assessment – A Brief History, Current and Extended Use for Security

In 1974, Norm Rasmussen from the Massachusetts Institute of Technology led a team from the Atomic Energy Commission to conduct the reactor *safety* study (WASH-1400) [5] in which they developed the concept of societal risk. The WASH-1400 study was published in 1975, and although widely criticized, it nonetheless established the foundational principles of probabilistic risk assessment (PRA) still widely used today. Shortly thereafter, a modified version of the societal risk model was first proposed for nuclear *safeguards (security)* [6]. Known as the ERDA-7 proposal, this approach was evaluated by Rasmussen, who concluded that safeguards (security) risk could not be quantified using the WASH-1400 developed societal risk approach [7]. Rasmussen said that he did not believe that risks involving malevolent human action could be quantified by traditional risk assessment methods like fault tree and event tree analysis because attack probability estimates could not meet important statistical requirements [7]. Over the years, the ERDA-7 proposal has been subject to reintroduction and modification [1, 8, 9, 10]. Similar to Rasmussen's conclusions, subsequent critical reviews stated an approach like ERDA-7 proposal based on traditional risk assessment not be used for security risk [11, 12]. The ERDA-7 approach is problematic for intentional malevolent acts, the terms in the equation are interdependent, data is, lacking which results in large uncertainties. Instead of using of the ERDA-7 approach, performance-based standards for the effectiveness of security systems as well as addressing consequences were recommended as useful tools [7, 13].

2.1. Current Definition of Risk

Kaplan and Garrick [14] stated the definition of risk that is most commonly used among modern risk analysts as, "Fundamentally... a risk analysis consists of an answer to the following three questions: *(1) What can happen? (2) How likely is it that [it] will happen? and (3) If it does happen, what are the consequences?* To answer these questions we would make a list of outcomes or 'scenarios' [where each line in the list] can be thought of as a triplet $<s_i, p_i, c_i>$ where s_i is a scenario identification or description; p_i is the probability of that scenario; and c_i is the consequence or evaluation measure of that scenario, i.e., the measure of damage. If this table contains all the scenarios we can think of, we can then say that it (the table) is the answer to the questions and therefore is the risk."‡ Thus, risk is defined as a collection of such triples, and since each scenario is

‡ Prior to Kaplan and Garrick, the most common definition of risk related to loss expectancy. Risk was defined as "probability *times* consequence." Kaplan and Garrick assert that risk is really "probability *and* consequence."

associated with a probability, one can summarize this set of triples as a "risk curve" which satisfies the definition of a statistical complementary cumulative distribution function (CCDF).

2.2. Estimating Security Risk Using the Current Definition

Security risk is frequently thought of in terms of three fundamental components: threat, vulnerability, and consequence [14]. These components are mapped into the above risk definition as follows. A scenario s_i represents a specific threat exploiting particular vulnerabilities to produce consequences. The likelihood of the scenario is composed of two parts: (i) the likelihood that the threat T with particular characteristics (e.g., number of attackers, weapons, tools, etc.) will attempt an attack (P_T), and (ii) the conditional likelihood that the attack by this threat will be successful ($P_{S|T} = 1 - P_{E|T}$). The consequences of a successful attack are represented by c_i. For many high-consequence facilities, attacks are so rare that statistical estimates of P_T are highly uncertain. As a result, analysts often neglect P_T and assess *conditional* risk, i.e., the risk that would exist given that the attack were to occur, on the basis of $P_{S|T}$ or $P_{E|T}$, or, for a DBT, $P_{E|DBT}$ [1]. When analysts assess threats, vulnerabilities and consequences, evaluating a range of possible attack paths, security risk is characterized by a set of conditional risk triples that exhibit many of the characteristics described above. These conditional risk triples are used to evaluate the efficacy of proposed risk mitigation options based on the degree to which they improve $P_{E|T}$ or $P_{E|DBT}$ for one or more scenarios s_i.

Conditional risk triples have a key drawback that limits their use in cost-benefit analyses: aggregated security risk cannot be computed because *conditional* probabilities $P_{S|T}$ or $P_{E|T}$ cannot be aggregated as a CCDF [7, 15, 16]. In order to perform this aggregation, several conditions must be met: P_T must be estimated, and the scenarios s_i must be mutually exclusive and statistically independent. In the world of physical security, this condition is clearly not met because intelligent and malevolent adversaries choose among scenarios and select the one that they believe to be in their best interest. In fact, adversaries even choose among scenarios that are not represented in a facility's scenario set because they may be deterred from attacking a target at one facility and choose to attack a different facility altogether. Consequently, P_T for s_i can never satisfy the necessary mathematical conditions required for aggregation as required in the traditional definition of risk [7].

Practical problems also exist when using the traditional definition of risk in a security context. First, P_T can only be estimated in a Bayesian sense and is enormously uncertain because we cannot know the intentions of all adversary groups. Historical attacks indicate that adversary choices are not random. Instead, adversaries assemble resources that they believe are sufficient to ensure a high likelihood of a successful attack, or they select targets and plan attacks that they believe they can successfully achieve within their available resources and abilities to execute. Hence, even a Bayesian estimate of P_T depends strongly on unquantifiable factors like dissuasion, deterrence and the adversary's level of goal commitment. Furthermore, P_T can change wildly over time as adversary groups are influenced by local and global political and social events of which we may not even be aware. Thus, the uncertainties in P_T are very large and can span several orders of magnitude for extreme but very rare attacks. Hence, investment decisions that are based on such risk estimates often cannot be supported with reasonable statistical confidence. Ironically, these uncertainties are caused in large part by the very definition of risk. Therefore, using attack probabilities for security risk is less useful than the comparable random event frequencies that make up safety risk analyses. This issue is compounded by the common representation of security risk not as a triple, but as a value obtained by multiplying P_T, $P_{S|T}$, and a metric representing severity of consequence that can also span multiple orders of magnitude.

In current security risk studies, when the collection of conditional risk triples is used to evaluate the efficacy of proposed risk mitigation options, the evaluations often rely on two key assumptions: first, the adversary embodies the DBT, and second, the adversary knows and exploits the scenario s_i with the highest likelihood of success. As a result, mitigation is effective to the degree that it increases the minimum value of $P_{E|DBT}$ across all scenarios s_i. Recall that in assessing $P_{E|T}$, one must assume that a specific attack scenario against a given target is carried out by an adversary with particular characteristics (e.g., number of attackers, weapons, tools, etc.). Thus, $P_{E|T}$ can change dramatically as different adversary characteristics are considered. Even small changes in the adversary characteristics required to defeat the security system can profoundly affect $P_{E|DBT}$ when the system's breaking point is near the DBT, but very large changes in the adversary characteristics required to defeat the security system can have a negligible effect on $P_{E|DBT}$ when the system's breaking point is not near the DBT. These situations can occur when the DBT is changed as new information indicates that adversary capabilities are increasing or as adversary behaviors reveal that prior threat assumptions no longer hold. The highly nonlinear relationship between the adversary characteristics required to defeat a security system and $P_{E|DBT}$ can make it difficult for security decision makers to prioritize security investment options on the basis of $P_{E|DBT}$ – especially when investments must be prioritized across multiple targets or facilities [17]. It can also result in security systems for which performance and costs are highly sensitive to changes in adversary capabilities and/or threat assumptions.

Using conditional risk for security assessment can also lead to an important unintended side effect. By focusing on the adversary's successes and failures during the hypothesized attack, the analyst can be led to focus only on security risk mitigation options that make the observed adversary successes less likely. In so doing, the analyst may not recognize risk mitigation opportunities outside of the actual attack execution. For example, it may be possible to deny the adversary certainty of information that is critical to attack planning, or to minimize the consequences of the attack through resiliency and redundancy. A holistic perspective is required to ensure that the most cost-effective security mitigation options are discovered and pursued.

2.3. Extending the Definition of Risk

To overcome the obstacles related to the use of probabilities with malevolent adversaries, we propose a modified definition of risk where, instead of considering the highly uncertain likelihood or probability of an attack, one considers its difficulty for an adversary to successfully accomplish against the target(s) under consideration. Thus, a security risk analysis consists of answers to the following three revised questions: *(1) What can happen? (2) How likely is it that [it] will happen? and (3) If it does happen, what are the consequences?* The triplet for security risk then becomes $<s_i, d_i, c_i>$ where d_i is the degree of difficulty for an adversary to successfully accomplish attack scenario s_i at a specific target in order to cause consequence c_i.[§] This definition explicitly acknowledges the observed adversary attack planning behaviors described above and addresses the problems associated with using probabilities to describe the intentional actions of both known and unknown intelligent actors. Risk evaluations using this definition do not require revision as adversary motivations change because this risk definition characterizes scenarios and targets rather than estimating the adversary's probability of attack. For each target, a number of scenarios can be posed, each correlating to a risk triplet. For a given consequence, there is a "threshold threat" that is the lowest difficulty (highest risk) scenario for an adversary to be successful.

2.4. Applications of Proposed Risk Definition to Security Risk

This work uses the proposed definition by focusing on estimating the minimum threat capabilities [or "threshold threat" (*TT*) characteristics] and degree of difficulty required for an adversary to accomplish a specific attack scenario that exploits a target's vulnerabilities and induces specific consequences with a reasonably high likelihood of adversary success $P_{S|TT} = 1 - P_{E|TT}$. Adversary attack preparation activities are viewed as a project planning exercise, wherein a planner has success criteria (e.g., adequate or desired consequences and thresholds for likelihood of success), and chooses among alternative strategies that meet these criteria (e.g., achievable resources and plausible attack

[§] This definition of risk, and specifically d_i, is a characteristic of scenario s_i for the specific *target*. The reader should not assume that d_i characterizes any specific adversary group or DBT. Rather, d_i incorporates the threshold threat characteristics needed for an adversary to have a high likelihood of success (i.e., a low value of $P_{E|TT}$) when attempting to execute scenario s_i at the specific target. It also incorporates the characteristics and complexities of the scenario that might make the scenario difficult for an adversary to accomplish successfully even if they had the requisite threshold threat characteristics.

scenarios), while considering the degree of difficulty that will be encountered in order to achieve a successful outcome. Investments reduce security risk as they either (a) increase the difficulty for an adversary to successfully execute the most advantageous attack scenario, or (b) reduce the severity of the scenario's expected consequences. The latter can be measured through existing consequence metrics, but measuring the former requires development of a reasonable and robust metric to characterize the adversary's degree of difficulty in achieving a "successful" attack with likelihood $P_{S|TT} = 1 - P_{E|TT}$. Thus, the proposed definition and metric build upon the well-known P_E-based assessment and design methods, but do not exhibit the strong nonlinear behavior that has been observed for $P_{E|DBT}$. Building this metric is not straightforward, as it requires one to compare and aggregate the relative degree of difficulty for disparate adversaries to successfully prepare for (e.g., acquire the requisite resources) and execute an attack (employ those resources in specific ways against specific targets). However, with such a metric, this definition of risk can form the basis of an objective risk-informed security analysis method. The proposed metric is described in Section 3.

Using the metric as a measure of scenario difficulty), an analyst can compare security risks by comparing attack scenarios' levels of difficulty and consequences. The insights from such comparisons can provide important and useful security risk management insights for a broad range of applications. The objective of a security decision maker might be thought of as follows: to make the easiest attack path as difficult as possible within the constraints imposed by cost, operational and programmatic considerations. Consider a decision maker who is responsible for several sites where each attack leads to similar consequences. Figure 1a shows how results from this method can be applied to security decision making. Each light-colored bar represents the difficulty of the *easiest attack scenario* at a notional site in its original (2007) configuration. Note how it was much easier for an adversary to achieve a successful attack at Site D than at any other site. Note also how security at Site B was already significantly better than the original (2008) goal level. The decision maker focused on improving security at Site D, and in 2010, security is much more balanced across the enterprise as the difficulty of the easiest attack is now roughly comparable across all sites (the top of the dark bar in the graph). The decision maker can justify to the funding source *why* particular security investments were made and describe the specific benefits that the investments produced. Further, if policy changes cause the security goal to change, the decision maker can explain in simple terms to the funding source why additional security investments are necessary. Prioritizing investments is straightforward for this application, and the method is compatible with computerized optimization programs.

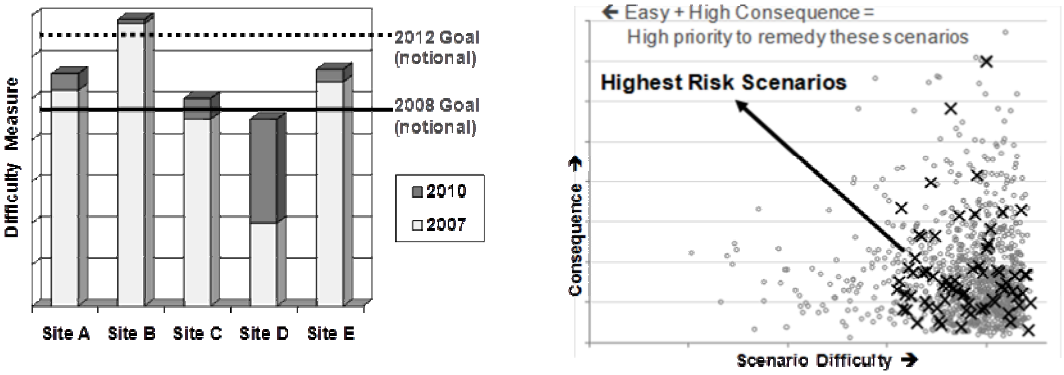

FIG. 1a (left) and 1b (right). (a) Comparing the relative difficulty of the easiest attack scenarios at five notional facilities where each attack leads to similar consequences. (b) Relative difficulty and consequences of attack scenarios at a notional facility (X symbols) compared with scenarios at other facilities within the enterprise (circles).

The situation where a variety of consequences are possible at a facility (or within an enterprise) is shown in Figure 1b. Here each identified attack path or scenario is represented as a circle on the scatter plot, with coordinates that represent the scenario's difficulty d_i and consequences c_i. Scenarios that produce higher consequences *and* are easier to accomplish are more attractive to an adversary because they represent a more efficient use of resources. Thus, they pose a greater risk and should be a higher priority for remediation. A scenario's risk can be reduced by reducing its consequence potential (moving the circle down), increasing its difficulty (moving it to the right), or a combination of these actions. Note that if one reduces the risk of a scenario s_j that is near the center of the pack of circles without also addressing scenarios that are more attractive (those that produce greater consequences *and* are easier to accomplish, *i.e.*, scenarios whose circles are above and to the left of s_j), the overall security risk may be unaffected by the investment because the most attractive scenarios remain available for adversary exploitation. Thus, the security investments should generally address those scenarios that are non-dominated (i.e., that represent the easiest way to produce consequences greater than or equal to c_j).

For Figure 1b from the perspective of the security decision maker for an enterprise, the X symbols represent the attack scenarios available at one facility, and that facility's manager wishes to mitigate the scenarios that are most attractive at that facility. The enterprise decision maker might use this graph, with circles representing attack scenarios available at other facilities within the enterprise, to inform the facility manager that only minimal security improvements will be supported because the enterprise has greater security risks that must be addressed first. On the other hand, if it is known from other sources that the facility is specifically targeted by credible threats, the enterprise decision maker may decide to support security upgrades at the facility anyway, believing that the easiest attack is not yet difficult enough.

3.0 Risk-Informed Management of Enterprise Security (RIMES)

The method described above is the Risk-Informed Management of Enterprise Security (RIMES). The following sections describe the general characteristics of the method and a system of metrics designed to describe and summarize the levels of difficulty that adversaries would face in successfully executing attack scenarios.

3.1. General Characteristics of the RIMES Method

The RIMES approach starts by identifying a scenario that would offer an adversary a reasonable expectation of success[**] against the target(s) under consideration, i.e., a scenario for which the conditional likelihood that the attack by this threat will be successful ($P_{S|T} = 1 - P_{E|T}$) exceeds a threshold established for this purpose. Such scenarios can be developed by any number of currently available means that are commonly used by the security analysis and vulnerability assessment community. Specific to each scenario, either explicitly or implicitly, are the resources (personnel, materiel, and knowledge) that an adversary would need to have, and the manner in which they would need to be employed, in order for the adversary to have a reasonable likelihood of success $P_{S|T}$ when executing the scenario against the target(s) under consideration.

Considerations of the difficulty for an adversary to mount this scenario are partitioned into the two essential phases of adversary efforts for any attack scenario - Preparation and Execution. Since adversary success in the scenario requires successful completion of both phases, they are viewed with comparable significance. The primary factors that are generally key to adversary success in each phase of attack have been identified through discussions with subject matter experts, review of various ranking schemes for adversaries or threats or scenarios, and analysis of a diverse set of specific scenarios. Since we require a metric that characterizes the relative difficulty of successfully (inducing and) exploiting target vulnerabilities, we express scenario success factors in terms of their manifestation at the interface between target and threat. For example, while level of funding can be important to adversary success, this is manifested at the target in other factors, such as quality and size of the toolkit used in the scenario. We have developed these factors so that they can be considered as roughly independent dimensions of generally equivalent importance.

In addition to reflecting key factors for scenario success, the required metric must also reflect the relative level of difficulty for adversaries to be successful in the scenario against the target(s) under

[**] For most attack scenarios, "success" means inducing a specific consequence of the adversary's choosing from the target.

consideration. To do this, five discrete levels of difficulty have been defined for each success factor dimension. Guidelines are being developed for analysts to consistently assign the appropriate levels to each success factor dimension in order to reflect the relative difficulty that an adversary would encounter to successfully achieve or acquire the characteristics required in that dimension for the scenario to succeed. It is important to note that this process does not assign adversaries to a particular level, nor imply that all dimensions of a scenario are at the same level. Rather, the process dissects a successful scenario into the minimum levels of difficulty associated with each of the key factors that generally underlie adversary success. Since the scenario is specific to the target(s) under consideration, this process characterizes targets in terms of the levels of adversary difficulty to recognize, induce, and exploit vulnerabilities that enable scenario success.

The levels of difficulty for the dimensions have been calibrated so that a particular level for one dimension roughly correlates to an equivalent level of difficulty for any other dimension. In general, the levels of difficulty correlate with the size of the portion of the spectrum of generalized potential adversaries that could reasonably expect to achieve or acquire the associated level characteristics. Level 1 characteristics are easily accessible or achievable by the general population, while Level 5 characteristics would typically be accessible or achievable only by elite forces or state supported operations. Different levels of difficulty are distinguished by different levels of costs, quality of leadership, law enforcement or intelligence signatures, time to achieve, availability, ingenuity, and/or sophistication.

3.2 Dimensions of Success for Attack Preparation and Attack Execution

As a basis for the difficulty of attack metric, Table I presents the dimensions of success for preparation and execution of adversary attacks. The dominant challenges for adversaries in the Preparation phase of efforts are in developing, acquiring, and preparing the resources – personnel, materiel, and knowledge - required for the scenario without being detected or interdicted by authorities. The dominant resource attributes that are keys to scenario success, and the primary considerations that differentiate levels of difficulty for the adversary to succeed, are described in Column 1. In the Execution phase, the manner in which adversaries employ their resources can also be critically important to their ability to succeed. The dominant success factor dimensions for attack execution, and the primary considerations that differentiate levels of difficulty for the adversary to succeed, are described in Column 2.

3.3 Calculating the Metric

Generalized guidelines (not presented here) have been developed for assigning one of five levels of difficulty to each of the attack Preparation and Execution dimensions for any particular scenario and target(s). A scenario for which an adversary is considered to have a reasonable expectation of success against the target(s) under consideration, i.e., for which ($P_{S|T} = 1 - P_{E|T}$) exceeds some threshold established for this purpose, is evaluated according to these guidelines. A numerical value is associated with each of the five levels of difficulty (currently, these are integer values 0 to 4). A dimension's values could also be weighted to reflect that dimension's relative general significance to adversary success, although research to date has not indicated a rationale for other than uniform weighting. Since the dimensions are roughly independent and span the most significant challenges that are key to adversary success, the level of difficulty for each of the phases of the scenario is calculated as the length of the vector described by the values along each of the phase's dimensions (an L_2 norm). Similarly, the metric for overall difficulty of that scenario for the target(s) under consideration is calculated as the length of the vector described by the levels of difficulty for each phase of adversary activity. This metric is specific to the scenario and target(s) under consideration.

4.0 Application of RIMES for Nuclear Facilities

RIMES has been applied to evaluate the theft and sabotage risks for two types of nuclear fuel cycle facilities – used nuclear fuel (UNF) storage and small modular reactors (SMRs).

Table I. Dimensions of Difficulty for Attack Preparation and Execution

Attack Preparation	Attack Execution
Active Outsiders: # of Fully Engaged Participants: the difficulty an adversary faces to successfully muster and prepare team(s) without alerting authorities, which increases with the number of participants.	*Ingenuity / Inventiveness:* the degree to which an adversary must be creative or ingenious in order to discover and/or induce, and exploit the vulnerabilities required for a successful attack. Low levels are associated with simple, straightforward attacks that can easily conceived by most adversaries, while high levels are associated with attacks that reflect unique, imaginative approaches that are more likely to surprise and befuddle even very well prepared defenses.
Active Outsiders: Training & Expertise of Fully Engaged Participants: the depth and diversity of expertise required of participants, and by the rehearsal required for tasks.	*Situational Understanding & Exploitation:* the level of acuity required by the adversary to recognize the occurrence of exploitable conditions and the flexibility required to leverage those opportunities. Levels of difficulty are differentiated by the transience, unpredictability and observability of vulnerabilities upon which success of the scenario depends.
Support Structure: Size, Complexity, and Commitment: the contributions required of a support base during attack preparation, e.g., intelligence, safe haven, training or staging facilities, finances, scientific or technological R&D, and manufacturing. Difficulty varies with the extent, diversity, and quality of contributions required, and the degree of engagement and awareness of purpose for these contributions.	*Stealth & Covertness:* the degree to which scenario success depends on the concealment or masking of attack execution activities in order to delay the point of initial detection and recognition by authorities. Levels of difficulty are differentiated by the existence, duration and multiplicity of undetected adversary operations that must be conducted within the observational purview of authorities.
Tools: Availability reflects the difficulty associated with acquiring the tools required to successfully execute a scenario. Tools can include weapons, transportation, breaching equipment, electronics, fixtures, armor, disguise, etc. The levels of difficulty are distinguished by factors that influence their availability: rarity, law enforcement / intelligence signatures associated with their acquisition or staging, and level of controls in place to protect against illicit usage.	*Outsiders: Dedication / Persistence / Commitment:* the significance of consequences at risk for the attackers, their support base, and/or their cause, the persistence of their risk exposure, and the degree of adversary certainty of those consequences.
Insiders: # of Contributors: one of three dimensions (key factors for adversary success) associated with contributions from insiders. Difficulty varies with the necessity for insider contributions, the number of contributors required, and the necessity of collaboration among multiple insiders.	*Insiders: Degree of Engagement & Risk:* the equivalent significance, persistence, and certainty of risk exposure required of insiders contributing to the attack.
Insiders: Security Controls on Contributors: contributions required from insiders that have greater levels of access to security-sensitive features are generally more difficult for adversaries to confidently acquire due to the security controls in place to mitigate the potential for such occurrences.	*Operational Composition / Complexity:* the required number, modalities, and orchestration of separate avenues of adversary attack execution operations. Modalities refer to the nature of vulnerabilities and exploitation operations required for the scenario: e.g., physical, cyber, procedural, etc.

4.1 RIMES for Used Nuclear Fuel Storage Security

For UNF, increased emphasis is being placed on extended storage, especially dry storage, potentially for many decades. As part of this emphasis, technical analyses and guidance documents are needed to assure that the security risks associated with extended storage are understood and minimized. Any assessment of security over a very long timeframe is a challenge. The security assessment needs to consider protection provided by a storage container (cask) as well as the facility protection measures and to address identified security issues over the timeframe of extended storage. RIMES is being applied to provide a framework within which to evaluate security risks that may change and evolve over the timeframe of extended storage. Attack scenarios have been developed for sabotage and theft and the difficulty of these scenarios evaluated. In general, the relative difficulty of attack for sabotage was moderate to high and for theft was very high. Evaluation of consequences, in terms of potential radiological releases, will be incorporated in future analyses. Additional scenario development will also consider changes in future conditions and alternative storage facility design concepts.

4.2 RIMES for Small Modular Reactor Security

A generic integral pressurized water reactor (iPWR) design [18] was developed to provide a basis for the RIMES analysis. This design pulled from many of the common features of iPWR designs currently available today without representing any one specific design. The work for SMRs identified a preliminary list of safety and support systems necessary for safe shutdown and then applied RIMES for example theft and sabotage scenarios. A total of 14 scenarios were evaluated to cover a range of attack types and consequences. Consequences were loosely binned into economic damage only, economic damage with release, core melt with little/no release, and core melt with release. Both outsider and insider attack scenarios were considered. Subject matter experts in reactor design, reactor safety, physical security, and response forces participated in the assessment. A long-term goal is to use these results to better inform physical security system design for plant designers. In many cases, rather simple design changes can either significantly increase the difficulty or reduce the consequence of a particular scenario.

In general, core melt (high consequence) scenarios were found to result in high difficulty levels. Multiple systems would need to be disabled, some of which are redundant. One scenario, which was found to be at a more moderate difficulty rating, could easily be remedied with a simple design change to the reactor building. Lower consequence property damage scenarios can be achieved with relative low difficulty—these scenarios do not lead to core melt or any release, but could cost the operator a significant amount of money in lost operational time. Scenarios with lower levels of difficulty can be addressed through design changes or improvements to the physical protection system that increase difficulty or mitigate consequences. The RIMES methodology made it much easier to examine cost-effective design changes. However, it should be noted that all of these scenarios will change when applied to specific vendor designs.

5.0 Conclusions

The RIMES methodology has been developed to address some of the key issues associated with applying traditional risk analysis to security. RIMES is an objective risk-informed method that is based on characterizing targets in terms of an adversary's degree of difficulty to prepare for and execute successful attacks. A focus on the level of difficulty of a particular attack as opposed to the probability of attack will enable decision makers to balance competing security interests (e.g., multiple facilities) and provide objective and unbiased justification for investment decisions, resulting in more robust and cost-effective security systems. This shift allows for designers to manage risk better by balancing increased security against those threats that require lower difficulty for an adversary to produce higher consequences.

This work has provided a preliminary examination of attack scenarios for two types of nuclear fuel cycle facilities – UNF storage and SMRs. For an individual facility, the RIMES methodology helps designers to focus on the attack scenarios of concern and the threats that can accomplish those scenarios, but RIMES can also examine how those scenarios and threats compare to those that could be executed in other parts of the nuclear fuel cycle. This work has also investigated and demonstrated how lower difficulty attacks for consequences of concern can be addressed by facility or security designs changes that can eliminate or mitigate the consequences or increase the difficulty of attack. The longer-term vision is to apply RIMES across the fuel cycle to examine most likely attack scenarios across various facility types to target investments to address security risks where they are needed most.

REFERENCES

[1] GARCIA, M.L., The Design and Evaluation of Physical Protection Systems, Second Edition, Butterworth-Heinemann (Elsevier), Burlington MA (2008).

[2] MATTER, J.C., Demonstration of Analytic System and Software for Evaluating Safeguards and Security (ASSESS), Proc. 30th Ann. Mtg. of Inst. of Nucl. Matls. Mgmt., Deerfield IL (1989).

[3] WALTER, K., Simulated Rehearsal for Battle: Joint Combat and Tactical Simulation, Sci. & Tech. Rev., Lawrence Livermore National Laboratory, Livermore CA, Apr/May (2009), pp. 16-22.

[4] Sandia National Laboratories, "Sandia National Labs' Security Risk Assessment Methodologies," viewed March (2010) www.sandia.gov/RAM/RAM%20Overview%20%20Presentation%20Aug%2006.pdf

[5] U.S. NUCLEAR REGULATORY COMMISSION, WASH-1400 Reactor Safety Study: An Assessment of Accidental Risks in U.S. Commercial Nuclear Power Plants, NUREG-75/014, U.S. Government Printing Office, Washington DC (1975).

[6] MURPHEY, W.M., SHERR, T. S., BENNETT, C.A., Societal Risk Approach to Safeguards Design and Evaluation, ERDA-7, Energy Research and Development Administration, Washington DC (1975).

[7] RASMUSSEN, N., Probabilistic Risk Analysis – Its Possible Use in Safeguards Problems, Proc. 17th Ann. Mtg. of Inst. of Nucl. Matls. Mgmt., Deerfield IL (1976).

[8] UDELL, C.J., CARLSON, R.L., Risk Evaluation System for Facility Safeguards and Security Planning, Proc. 30th Ann. Mtg. of Inst. of Nucl. Matls. Mgmt., Deerfield IL (1989).

[9] UDELL, C.J., et al., Short Form Risk Evaluation Method, Proc. 34th Ann. Mtg. of Inst. of Nucl. Matls. Mgmt., Deerfield IL (1993).

[10] BIRINGER, B.E., et al., Security Risk Assessment and Management – A Professional Practice Guide for Protecting Buildings and Infrastructures, John Wiley & Sons, Inc., Hoboken NJ (2007).

[11] RICHARDSON, J.M., Comprehensive Safeguards Evaluation Methods and Societal Risk Analysis, SAND82-0366, Sandia National Laboratories, Albuquerque NM (1982).

[12] COX, JR., L.A., Some Limitiations of Risk = Threat × Vulnerability × Consequence for Risk Analysis of Terrorist Attacks, Risk Analysis, **28** (2008) No.6.

[10] SNELL, M.K., GARDNER, B.H., Determining System Effectiveness Against Outsiders using ASSESS, Proc. 32nd Ann. Mtg. of Inst. of Nucl. Matls. Mgmt., Deerfield IL (1991).

[13] KAPLAN, S., GARRICK, B.J., On the Quantitative Definition of Risk, Risk Analysis, **1**:1, (1981).

[14] U.S. DEPARTMENT OF HOMELAND SECURITY RISK STEERING COMMITTEE, Risk Lexicon, U.S. Department of Homeland Security, Washington, DC (2008).

[15] WYSS, G.D., et al., Risk-Based Cost-Benefit Analysis for Security Assessment Problems, Proc. 51st Ann. Mtg. of Inst. of Nucl. Matls. Mgmt., Deerfield IL (2010).

[16] WYSS, G.D., et al., Risk-Based Cost-Benefit Analysis for Security Assessment Problems, Proc. 44th Ann. Intl. Carnahan Conf. on Security Tech., San Jose CA (2010).

[17] WYSS, G.D., et al., Risk-Based Cost-Benefit Analysis for Security Assessment Problems, Proc. 50th Ann. Mtg. of Inst. of Nucl. Matls. Mgmt., Deerfield IL (2009).

[18] LEWIS, T., et al., Generic Small Modular Reactor Plant Design, SAND2013-10406, Sandia National Laboratories, Albuquerque, NM (2012).