POINT PAPER ON NUCLEAR MATERIALS SECURITY

Draft – August 26, 2002 – Draft

Introduction

As part of the Nuclear Regulatory Commission's (NRC's) comprehensive review of its safeguards and security program, NRC staff in the Offices of Nuclear Security and Incident Response, Nuclear Material Safety and Safeguards, Nuclear Regulatory Research, and State and Tribal Programs has been assessing the risks of malevolent attack on materials facilities and activities. The purpose of this paper, after it has been reviewed by senior NRC management, is to provide a consistent conceptual approach for proceeding with these assessments, including objectives, consequence criteria, assumptions, and general logic.

Security Risk

The risk of a hostile or malevolent act involving radioactive material (e.g., radiological sabotage, deployment of a Radiological Dispersal Device or Exposure Device, theft of strategic special nuclear material) can be defined as follows:

$$\text{Risk} = P_{attack} \times (1 - P_{interrupt}) \times \text{Consequences}, \text{ where}$$

P_{attack} = Probability of attack or malevolent act (initiating event),

$P_{interrupt}$ = Probability of interruption or defeat/disruption of attack, and

Consequences = Adverse impacts to be avoided, such as radiation injuries, loss of human life, economic impacts, or societal disruption.

Consistent application of this equation allows the NRC to employ a coherent approach to ensuring the protection of the public across the broad spectrum of facilities and activities regulated under the materials program. The level of risk should be sufficiently small to ensure high assurance of protection of the public from malevolent attacks or uses of byproduct, source, and special nuclear material. A quantitative risk objective could be established analogous to the quantitative objectives in the NRC's Safety Goal Policy. However, as discussed later, it may be difficult to quantify the level of security risk associated with licensed facilities and activities.

Consequences

The simplicity and direct proportionality of the risk equation makes the proper and defensible selection of input parameters all the more important. The

"Consequences" term needs to be further defined in terms of the adverse impacts to be avoided. For the purposes of this analysis, these adverse consequences include the following:

- A public fatality, outside of the authorized location of use for the radioactive material, caused by acute radiation or chemical exposure;
- Economic impacts exceeding $50M due to loss of service, decontamination and waste disposal, and related costs; or
- Societal impacts sufficiently large that would significantly undermine the confidence of the public in the ability of the government (local, State, or Federal) to fulfill the functions for which that government was established.

These consequences then become the fundamental objectives of the security program. In other words, sufficient security is needed to guard against (1) one public fatality from radiation or chemical exposure, (2) more than $50M in economic damages, or (3) significant societal impact. If these objectives can be satisfied without additional security, then no security enhancements are necessary. On the other hand, if these objectives can be violated under existing conditions, then additional security enhancements may be appropriate and necessary to ensure adequate protection of the public.

If it is inconceivable that any kind of attack could cause consequences that exceeds these criteria, then the maximum consequences of any hostile attack equal zero and there is "no risk." This is the simplest solution to the risk equation because it avoids the bother of attempting to define the other terms. In addition, in such cases, no additional effort is warranted to enhance security because no value will be gained by such efforts.

Probability of Attack

For situations in which consequences cannot be zeroed out, the NRC approach continues with the estimation of the probability of attack. The probability of attack has traditionally been difficult to define or quantify in security programs. Unlike accidents initiated by approximately random events (e.g., earthquakes, tornadoes), attacks and other malevolent acts are instigated by humans who generally observe their targets, plan the attacks, and wait for moments of opportunity to increase their chance of success. In addition, "attacks" can be carried out by different types of adversaries ranging from one or more terrorists, to petty criminals, organized crime, disgruntled employees, radical organizations, or extremist group or individual. Consequently, "attack" may include anything from full frontal assault intended to cause radiological or chemical sabotage, to theft of material, extortion, ransom, economic loss, cessation of offensive practices, or revenge. The probability of attack will also reflect the motivations, intent, and capabilities of the adversary.

The probability of attack may also vary as a function of the notoriety of the target. For example, certain nuclear facilities are known throughout the world because of the significance of events that occurred at these facilities. With greater prominence, it is reasonable to assume that a terrorist may place greater weight on attacks on such facilities because (1) they are generally better known to a greater number of people, (2) information about them is generally more available in news and cyber media, and (3) adverse consequences caused by these facilities would have a perceived greater impact.

Further, the probability of attack is a function of the relative ease of attack, which reflects the relative effectiveness of security measures in place to guard against the attack. Thus, the probability of attack for a facility with obviously robust security measures would be expected to be less than for a facility with lesser security, even though both may possess material of equal attractiveness and pose the same potential for inflicting consequences. The general assumption is that an adversary will be more likely to attack the facility with the weaker security to maximize the probability of success in the attack. This also introduces a potential dependency between the probabilities of attack and of interruption, for as the probability of interruption increases due to use of more effective security measures, so to would the probability of attack potentially decrease if the security measures are apparent and have a deterrent effect on the attack.

For all of these reasons, it is extremely difficult, if not impossible, to develop quantitative estimates of the probability of attack on a given nuclear facility or activity. The historical number and characterization of attacks on nuclear facilities or activities is insufficient to develop a quantitative estimate of the probabilities of attack. In addition, the Intelligence Community and the Law Enforcement Agencies are unable to estimate the likelihood of an attack due to imperfect knowledge of adversary capabilities, motivations, and intent. Consequently, the best that can be accomplished is a qualitative estimate of the likelihood of attack (e.g., very low, low, medium) based on consideration of the factors described above and on intelligence and law enforcement information. Although qualitative, this information can be useful to develop relative estimates of the security risks associated with attacks on nuclear facilities and activities.

Probability of Disruption

To quantify the probability of disruption, it is necessary to carefully assess the effectiveness of existing security measures in accomplishing the objectives described above. These measures would either deter the attack; prevent access to sensitive material, equipment, people, or information; defeat the attack (e.g., neutralizing the adversary); or delay the adversary long enough to allow the adversary's apprehension by law enforcement officials. If the collective set of security measures fails and the adversary achieves its objectives, then the probability of disruption goes to zero. In contrast, if there is no likelihood of the

adversary succeeding, then the probability of disruption goes to one, and the overall risk goes to zero.

Quantitative estimation of the probability of disruption requires computerized vulnerability analysis, where security officers simulate interdiction and defeat of simulated adversaries. These kinds of simulations require fairly extensive characterization of the facility, existing security barriers, intrusion detection and assessment, alarm systems, and response timelines. This level of characterization may exceed present capabilities at most, if not all, civilian nuclear facilities, especially at the smaller materials facilities, which have historically had minimal security measures and protection against malevolent acts.

However, semi-quantitative estimates can be developed using performance testing for security systems, such as rigorous force on force exercises, security tabletop drills, timeline analysis, and similar examinations. For extremely simple security systems, crude estimates of the probability of disruption can be developed using fault trees to examine the effectiveness of security measures in detecting, defeating, or delaying the adversary.

Design Threats

Some consideration of the threat is necessary to assess the probability of disruption. Obviously, a larger number of military-trained and dedicated adversaries may fare better in attempting to defeat security measures, than a less skilled, less-well equipped adversary with fewer numbers. Assessment of the current threat environment alone yields an incomplete basis upon which to estimate a design basis threat or even design characteristics for the threat. This again is a limitation of the imperfect knowledge by the Intelligence Community and law enforcement of the motivations, intents, and capabilities of adversaries. Although it is possible that the adversary would have a large number of heavily equipped, skilled, and dedicated attackers, it is less certain that the adversary would commit these capabilities in an attack on a nuclear facility or activity that would not or might not cause the impacts intended by the adversary. For example, it has traditionally been assumed by the U.S. Government that the threat against nuclear targets is proportional to the potential adverse consequences of an attack against the target. The largest threats have been postulated for nuclear facilities that possess fully functional nuclear weapons, whereas the smallest threats have been assumed for facilities and targets that pose far less significant consequences if the adversary succeeds. The Department of Energy is currently considering this approach in developing its revisions to the Design Basis Threats following the terrorist attacks on September 11, 2001.

The NRC approach maintains this same general concept by applying a graded design threat to civilian nuclear facilities and activities. The largest threats in

terms of adversary attributes (numbers of attackers, level of dedication, adversary skill, weapons and equipment) are assumed to apply to nuclear facilities that pose potential greater consequences, such as nuclear power reactors and Category I fuel fabrication facilities. Smaller threats are assumed to apply to other nuclear facilities and activities. The smallest threats would apply to most materials facilities, which are substantially less notorious and visible than the power reactors and fuel fabrication facilities. In addition, the potential consequences from such materials facilities are substantially less than the other nuclear targets. Moreover, the materials facilities are more similar to other industrial, academic, and medical facilities, which are not currently defended analogously to nuclear power reactors and fuel facilities. These latter facilities are uniquely secured among private sector facilities in the critical infrastructure.

Nevertheless, for the purpose of assessing the vulnerabilities of materials facilities to confirm that larger consequences are not feasible or to assess the probability of disruption, the threat characteristics used for power reactors and Category I fuel facilities provide an upper bound to the characteristics of the threat that might, although unlikely, be experienced at a materials facility or activity. There is no need to contemplate the potential consequences or vulnerabilities of materials facilities to threats that exceed the characteristics and attributes assumed to be active for nuclear power reactors or Category I fuel fabrication facilities and material. However, as described above, NRC will not necessarily assume threats up to this magnitude apply for materials facilities, unless intelligence indicates that this is the case or such assumptions are warranted because of the potential for significant consequences of such attacks.