D920612

The Honorable Ivan Selin
Chairman
U.S. Nuclear Regulatory Commission
Washington, D.C. 20555

Dear Chairman Selin:

SUBJECT: IMPLEMENTATION OF THE SAFETY GOAL POLICY

In our report of December 18, 1991, we expressed reservations about a staff proposal contained in SECY-91-270, "Interim Guidance on Staff Implementation of the Commission's Safety Goal Policy," and offered to produce an alternative approach after we had further considered the matter. You urged us to do so. During our 386th meeting, June 4-5, 1992, we developed the following proposal for consideration.

We assume that the Commission desires to establish guidance for the NRC staff to ensure that regulatory activities will be conducted in a manner consistent with the intent expressed in the Safety Goal Policy Statement. Although the plan outlined in SECY-91-270 purports to provide such guidance, we have major reservations about it. First, it applies only to a small part of the spectrum of regulatory activities and should not be characterized as the plan to implement the policy. A more strategic vision for implementation is needed. Second, even as a tactical tool, part of an overall implementation program, the SECY-91-270 plan has some significant inconsistencies with the policy.

We interpret the safety goals to be an expression of "how safe is safe enough." Thus, the Policy Statement expresses the Commission's intention that the safety of the general population of plants should be consistent with the goals, but implies no requirement or expectation that either individual plants or the population of plants must surpass the goals. The Safety Goal Policy Statement defines an acceptable level of safety for the nuclear enterprise.

This means that regulatory programs should not be an unending quest for higher and higher nominal levels of safety, but should be directed instead toward providing assurance that the plants, as a whole, meet the standard of safety already proclaimed by the Commission.

COMMISSION USE OF THE POLICY

The most important use of the safety goals should be by the Commission itself. The goals describe the level of safety which the Commission promises to achieve through its regulatory efforts. The body of regulatory activity in place is, in effect, its current effort to implement the safety goal policy. Whether this present activity is adequate can be partially evaluated by comparing the fruit of these practices, the safety level of operating plants, to

the goals.

Data on safety performance are available from two sources. The first is the set of "bottom-line" risk estimates from the many available PRAs, including those described in NUREG-1150 and those being developed under the Individual Plant Examination (IPE) program. Imperfections in this body of data are clear and should be recognized, but the data contain important information. One defect is incompleteness, e.g., while the PRAs characterize the design and physical status of plants fairly well, they say little about how well the plant is operated.

The second source is the risk information deduced through the Accident Sequence Precursor (ASP) program from operating experience. These data are also imperfect and must be judiciously used, but they also contain important information.

An indication from present PRAs is that the general population of plants operates near the standard set by the safety goals. One might conclude that the present level of regulation is sufficient. But uncertainties abound.

Assessment of the ASP data is also difficult. Improvements in ASP methods are needed before solid conclusions can be drawn from them.

No general conclusion is apparent at this time. We are not sure any can be drawn beyond a general impression that there are no indications that U.S. plants are failing to operate in accordance with the Commission's goals. However, we believe these (PRA and ASP) are important attempts, really the only quantitative attempts, to evaluate the overall safety performance of U.S. nuclear power plants and, presumably, the effectiveness of the Commission's present efforts to implement the safety goal policy. One cannot know whether changes in regulations are necessary and sufficient unless one has some measure of how effective the body of regulations has been in fostering a population of plants which operates in accordance with the safety goals.

STAFF USE OF THE POLICY

Although we consider the strategic use of the safety goals, as discussed above, to be the most important, it is also appropriate, when feasible, that the policy serve tactical purposes. It is a mistake to expect that a tactical use can be easily correlated with the high-level safety goals. For example, in a narrow sense the Backfit Rule (10 CFR 50.109) could be considered to be in conflict with the goals, if the latter are accepted as a statement of "how safe is safe enough." Also, one inevitably confronts the uneasy relationship between the safety goals and the definition of adequate protection. Rather, the principal use of the safety goals in support of tactical decision-making tools by the staff should be to help ensure conformance with the policy.

One tactical use would be to evaluate the need for proposed enhancements to regulations. Strictly speaking, we don't have much basis, from the perspective of PRA and ASP studies carried out to date, to argue that enhancements are needed. However, the

assessment of the safety of operating plants is an immature and incomplete undertaking. Until such assessments are further along, proposals for regulatory change, based on judgments about safety or on better understanding or knowledge, will evolve.

Given a particular contemplated regulatory action, a procedure can be developed to use surrogate guidelines, derived to be consistent with the safety goals, to decide its advisability. However, there should be no expectation that unambiguous yes-or-no answers can be established with such a procedure, simply by comparing risk estimates to the surrogate guidelines. Allowance for uncertainty and for unquantifiable factors must be included.

There are many proposals for additional requirements. These include resolution of matters typically brought before the Committee to Review Generic Requirements (CRGR). There is also systematic activity to help decide whether development of new requirements should be undertaken, e.g., the Generic Safety Issues (GSI) "prioritization" program. The proposal in SECY-91-270 was developed to evaluate the first of these. We understand that a similar proposal is being developed to provide guidance as to whether the initial development of new requirements should be undertaken. Neither of these proposals is inappropriate in concept; they could help ensure that these important staff activities reflect the intent of the safety goal policy. However, as we stated in our report of December 18, 1991, much of the SECY-91-270 plan seems to miss the full intent of the policy.

There are not many current proposals for deletion of requirements. This is unfortunate. There have been programs directed to this end, but they seem to have come up dry.

In our report dated May 13, 1987, on the Safety Goal Policy Statement, we commented on the use of surrogates to facilitate application of the safety goals to lower-level regulatory problems. In that report, we emphasized the importance of structuring the surrogates so that they remain surrogates, and do not become new de facto safety goals, more conservative than the original ones. While we emphasized the avoidance of excessive conservatism (because it is more of a problem at NRC), there are also pitfalls to be avoided in the other direction. It is Scylla and Charybdis, and one requires precise navigation to avoid the perils on either side.

But precise navigation is not possible in the world of PRA, the necessary tool for implementation of the safety goals. Any calculation of a probability has inevitably bound with it an uncertainty, and the uncertainty, however expressed, is as much a result of the analysis as the central number. We said more about this issue in our report dated December 14, 1991, on the use of PRA by the staff. (The undeniable allure of a precise decision mechanism leads all too often to staff decision making based on single bottom-line probability estimates, with at best lip service paid to uncertainty. We need to do better here.) While surrogate measures of risk at lower levels of aggregation can be invented and can be expressed as precise numbers, their tactical use must reflect the uncertainties associated with their calculation. It is

important to distinguish between the statement of a surrogate (or indeed a goal) as a precise number, and its calculation for some given situation, which will contain uncertainty. Without this distinction, irrational decisions are not only possible, but are certain. If a calculation which is uncertain by a factor of 10 shows that a proposed rule change exceeds some threshold criterion by 10 percent, is it rational to implement the change?

Given the first caveat about avoidance of added conservatism through surrogation, various lower-level surrogates for the safety goals have been suggested over the years, applicable in different situations. Among them are probabilities of 1E-6 per reactor-year for a large release, 1E-4 per reactor-year for significant core damage, and 1E-1 for a conditional containment failure probability. Each of these deserves continuing consideration to provide assurance that it is neither unduly conservative nor the converse, each can be calculated with substantial (and quantifiable) uncertainty, and each seems to us a reasonable step toward a useful surrogate for the full safety goals in a regulatory decision-making process.

What is still needed is a means for incorporating the necessary uncertainties in the calculations into the decision making. There will be some cases for which the calculated effect of a proposed change will be so clearly above or below the requirements of the surrogate standard (taking the uncertainty into account) that the decision process is simple and beyond reasonable disagreement. This might be judged by choosing some appropriate statistically described confidence level for the ordering of the surrogate standard and the calculated effect of the proposed change.

The difficult problems appear when there is inadequate statistical confidence that the proposed change meets the threshold surrogate standard; presumably this will be far from an uncommon event. For such situations we can only say that there is no free lunchýif the probabilistic situation is uncertain, other criteria will have to be used to bring the matter to a conclusion. This is not so strange; before the probabilistic era began, these other criteria were all that was used in decision making. In this proposal they would be used only to augment the safety-goal-based considerations, and then only in the event of substantial uncertainty.

What are these other criteria? Apart from shibboleths like engineering judgment, they include optimization of effort, resource allocation, discounting of impact timing, the intangible safety benefits of stability (if it ain't broke, don't fix it), the number of plants affected (there may well be times when a proposed change to a few plants will appear desirable from the point of view of the safety of those plants, but they are so few that the impact on the public risk will be small), and a host of other considerations. We leave their invention to the Commission and the staff, and wish only to note that there are times when decision making is difficult.

We do note, as a matter of principle, that there is no probability that cannot be quantifiedýthe only issue is the level of uncertainty associated with the quantification. By the same token,

there is no usefulness to a calculated probability without an associated statement, in some quantitative form, of its uncertainty.

We think the scheme we have outlined above is workable, though we recognize that we have provided only its skeleton. We also recognize that there will have to be a learning phase, in which the staff subjects proposed enhancements (and the converse) to the kind of analysis described here. Indeed, there will have to be a learning process for the more strategic implementation of the safety goals that we described in the first part of the letter. If the Commission subscribes to this general approach, we will be happy to work with both you and the staff to bring this long enterprise to a constructive conclusion. At best, it can offer a structure for more efficient use of NRC resources, and more effective regulation of industry, by focusing attention on regulatory activities calculated, however imperfectly, to have the most impact on the health and safety of the public.

Additional comments by ACRS Members Harold W. Lewis and J. Ernest Wilkins are presented below.

Sincerely,

David A. Ward
Chairman

Additional Comments by ACRS Members Harold W. Lewis and J. Ernest Wilkins

Although we thoroughly approve of the direction proposed in this report, we regret that the Committee has chosen to make the reference to confidence levels on page 4 of the report so terse that the recommendation conceals real problems.

The suggestion that a decision-making mechanism can be based on a confidence level is workable, but the choice of a specific levelý90% or 95% or whateverýinvolves a balance of benefits and effort that can only be resolved by the Commission. It is not a matter for fiat, but for analysis.

This is especially difficult for two reasons. The simplest one is the fact that the probability distributions in most PRAs are non-gaussian, so the familiar translation into a sigma level is inappropriate. That makes the application of a confidence criterion less straightforward than might appear on the surface, and departs from many engineers' experience.

Far more important, and ignored by the Committee, is the fact that the words "confidence level" mean different things to different people. Engineers tend to have little education in the subtleties, classical statisticians have little experience with low-probability analyses, and classical and Bayesian (we would say modern) statisticians both use the term "confidence level," but mean

entirely different things by the term. These are not just semantic differencesýthey need to be resolved if a confidence criterion is to be used, lest the ambiguities render a difficult job impossible.

References:

1. SECY-91-270 dated August 27, 1991, from James M. Taylor, NRC Executive Director for Operations, for the Commissioners, Subject: Interim Guidance on Staff Implementation of the Commission's Safety Goal Policy
2. Staff Requirements Memorandum dated February 21, 1992, from Samuel J. Chilk, Secretary, for James M. Taylor, NRC Executive Director for Operations, William C. Parler, General Counsel, and David A. Ward, Chairman, ACRS, Subject: SECY-91-270 - Implementation of the Safety Goal Policy Statement
3. Reports by the Advisory Committee on Reactor Safeguards on implementation of the safety goal policy:
   a. ACRS Comments on an Implementation Plan for the Safety Goal Policy, dated May 13, 1987
   b. SECY-91-270, Interim Guidance on Staff Implementation of the Commission's Safety Goal Policy, dated December 18, 1991
4. U.S. Nuclear Regulatory Commission, NUREG-1150, "Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants," December 1990