UNITED STATES
NUCLEAR REGULATORY COMMISSION
WASHINGTON, D. C. 20555

November 14, 1980

MEMORANDUM FOR: File

FROM: George Sege George Sege

SUBJECT: CRITERIA FOR SAFETY GOAL FORMULATION: A WORKING PAPER

CONTENTS

8012030 877

PURPOSE AND SCOPE

The purpose of this working paper is to set down, for preliminary discussion, potential criteria by which the merits of a safety-goal formulation should be judged. The paper represents an initial step in the development of such criteria, for use in policy development -- notably for near-term use in connection with the Preliminary Policy Paper that is being prepared as part of the Commission's program to develop a safety goal (NUREG-0735).

The paper attempts to set down a range of possible serviceable criteria and discusses available options, without proposing one definitive set of criteria at this time.

In preparation of this paper, information bearing on criteria from various sources (Decision Research, ACRS, IEEE, preliminary discussions with persons inside and outside NRC, etc.) was freely used -- and freely modified. The general approach and analysis are the author's own.

CATEGORIES OF CRITERIA

Criteria for judging the merits of a safety-goal formulation can be grouped into six categories: practicality, reach, result, clarity, public acceptance, and utility. Certainly, the categories overlap to some extent; in fact, the last category -- utility -- can be viewed largely in terms the other criteria as component elements. Nevertheless, it is suggested that the proposed taxonomy could be useful in analysis and discussion.

The discussion that follows will address each of the categories of criteria in turn.

Criteria of Practicality

Here the basic question is:

"To what extent is the formulation consistent with realistic recognition of the conditions and constraints under which it is generated and applied?"

Two specific criteria are of prime importance here:

. Recognition of political or ideological motives behind safety concerns; with appreciation of their nature and varieties, sponsorship, intensity, durability, and mode of manifestation.

. Recognition of uncertainties concerning probabilities and consequences of unwanted events, including uncertainties that can be narrowed by experience, research, and analysis, and those that cannot.

Also included among conditions referred to here are the following:

. Dichotomies between public and technical perceptions.

. The role of secondary consequences -- such as, for example, the impact of the TMI accident not only in terms of shutdown of TMI-2 but in its general effect on the nuclear power industry and its regulation.

Criteria of Reach

The reach of the safety-goal formulation involves two related basic questions:

. To what extent does the formulation address significant elements of the issue?

. How complete is it?

These criteria derive importance from the very nature of an agency safety-goal statement: such a statement would seek to rationalize safety regulation at a higher level of aggregation of issues than other regulatory instruments (rules, case decisions, etc.).

A part of these criteria relates to the generality of the policy; another, to what special issues are included, notably:

. Safety-cost tradeoffs

. Relation between goals for plants in operation, being built, or planned.

. Goals for future improvements.

There is a widely shared view that the core issue in safety-goal formulation relates to power-reactor accident considerations -- more particularly, major accidents. Specific criteria within this core issue could relate to whether and to what extent the stated goal policy includes within its sphere of determination or influence: rules and case decisions on plant design, operation, siting, emergency planning, degraded-core provisions, anticipated transients without scram, off-site contingencies, etc. Also, whether it addresses public impacts only or in-plant and precursor events too; whether it includes a policy on decisions under uncertainty; and the extent to which it excludes accident considerations that pose difficult policy problems (e.g., any implications for dilemmas of operational discipline vs. individual rights).

More broadly, at issue here is to what further extent (beyond major accidents) power-reactor safety is covered by a candidate policy approach: does it include risks of minor accidents and normal operation; sabotage; war? A case can be made here for giving particular weight to sabotage, because war is the province of other governemntal entities, and normal operation (except for its accident potentials) has had a safety record that suggests absence of major policy issues -- and is, further, subject to radiation-safety goals set by EPA.

A further broad criterion in this category is whether the goal covers -- or can readily be adopted to cover -- nonreactor facilities and processes that are inexorably entailed by reactor operation, or even regulatory areas not so entailed. The most notable issue here is that of radioactive wastes. Fuel reprocessing, should it be authorized, would figure here. Less prominent areas include other steps in the fuel cycle.

The degree to which the goal policy affects the decisions concerning the issues that it reaches is another criterion in this category.

Criteria of Result

At issue here is the extent to which the policy leads to coherent and rational results.

It is an underlying assumption in efforts to develop a safety goal that a cohesive rationale connecting the important interrelated issues in safety regulation can lead to better safety and better management of resources devoted to safety than disaggregated decisions alone.

A result-oriented test of a proposed policy is whether it can be applied to different aspects of an issue cluster and whether, when so applied, it leads to reasonable results in each. (E.g., siting, containment features, degraded core, ATWS, operator training).

These are the criteria that watch for pitfalls in terms of unintended effects. Specific criteria include such as these:

. Does the policy channel attention and resources to issues that are both significant and tractable?

. Does it leave room for initiatives to improve safety?

. While striving for coherence, does it allow enough flexibility for addressing related issues (especially in view of the uncertainties in them) differently when it is reasonable to do so. (A recently rected difficulty that illustrates this issue is the impact of the old LPZ concept on emergency planning requirements.)

Criteria of Clarity

The question of how clearly a safety goal can be interpreted in terms of regulatory actions (on which hinges the predictability of such actions) has two parts:

. How clear is the goal?

. How clearly can it be known whether a particular action serves or is consistent with that goal?

The first of these parts may be largely a matter of definition of terms. If the policy calls for reactors being "very safe", what does that mean? If it calls for a "major accident" having an annual probability of occurrence of 10^{-x}, what sort of accident is meant?

The second part of the question is primarily one of verification, of measurement. How can one know, or reasonably judge, whether some low probability target (say, 10^{-6} per reactor-year) is met? It should be noted that absence of strict verifiability, in a satistically meaningful sense, does not doom all quantitative goals for high-consequence, low-probability events to failure on this criterion. However, available options to deal with the verification problem may introduce weaknesses

with respect to other criteria. This is illustrated by reference to three possible approaches to the verification problem:

. Use of the top-level probability target can be limited to decisions concerning outliers: the clearly prohibitive and the clearly trivial. By leaving many decisions outside its purview, this approach detracts from the policy's value on the criterion of reach.

. The decisions as to whether a goal is met can be entrusted to a group of specially empowered experts. (We avoid the term "science court".) Here there is some inherent danger to the rationality of the regulatory results, since decisions would be made by methods that would necessarily in part transcend the scientific method. Predictability could also suffer, as experts and their opinions change.

. Verification could be tied to a prescribed mode of calculation. (E.g.,: "Consider the aggregated probability to be the sum of the probabilities of the ten dominant sequences described below.") With such approaches the quality of the results depends heavily on the fidelity of the calculational model; one needs to consider whether there is enough respect for uncertainties (including possibilities of crucial oversight), called for by the criterion of practicality.

A separate and further criterion in this category is that of clarity in terms of being understable to the general public, as well as the more limited publics directly involved, since broad public interests are affected.

Criteria of Public Acceptance

The question of the extent to which a safety goal is acceptable to affected publics is complicated by heterogeneity. The various publics (the industry, nearby residents, the general public and its various interest and opinion segments) are affected differently. Values and perceptions differ and there are differing presures to accommodate different views.

Ultimately it is the public in its social and political forms that determines the values that the safety goal must serve. The agency's goal is subject to judgment with respect to the sensitivity to public senses of value that it reflects.

There is a particular burden on the agency to seek to impart to the public a sound perception of risks addressed and of the relation between policy and result. A significant criterion in the public-acceptability category is whether a policy is distorted by catering to public misperceptions, or reflects courage and skill in trying to correct possible damaging errors in public understanding of the facts.

Criteria of Utility

The criterion of "how useful is it?" is not a substantially independent criterion. Rather, it is largely a resultant of the application of the other criteria discussed. Nevertheless, it merits separate consideration, because it involves a synthesizing judgment as to how the results of evaluation from other aspects are reflected in a degree of usefulness, as well as whether some special factor not foreseen in the other criteria might defeat a candidate approach's utility.

It is in this category that the question of proper balance among the several other criteria belongs. We have, for example, commented earlier in this discussion on possible conflicts between the demands of clarity and of quality of results.

Also involved here is the practical matter of applying the policy -- the extent to which the potentialities of the safety-goal policy can in fact be expected to be effectively translated into useful regulatory and licensing results by NRC staff, Licensing Boards, the industry, intervenors, the courts.

OPTIONS

In evaluating proposed safety-goal approaches, all six of the categories of criteria discussed are essential, in the limited sense that a severe enough shortcoming in any one of the categories could render a candidate approach unacceptable. Certainly (to account for each of the six criterion categories), no approach can be viable if it is divorced from its real-world context, reaches too narrow a slice of the issues, leads to unreasonable results, cannot be interpreted in terms of regulatory actions, is rejected by the body politic, or is useless.

However, a wide range of evaluation options exists in terms of the degree of emphasis accorded to the various categories of criteria, or to specific criteria within each category. Thus, one might emphasize comprehensiveness of the policy, at the cost of toleration of interpretive difficulties. Alternatively, clarity and ease of interpretation could be emphasized, at the cost of limiting the policy's reach to the more readily tractable parts of the issue.

On the basis of the foregoing discussion, specific sets of criteria, with assigned degrees of emphasis, could be structured into a selected number of discrete evaluation-method options. We do not do so at this time, and we leave open for now the question of whether such discrete options should be constructed later in the Commission's safety-goal program.

We do note, however, that in the program's early stages (keyed to the Preliminary Policy Paper) the criteria should be applied tolerantly, so that candidate approaches would not be too easily rejected from further consideration. As the program progress towards the narrowed set of options of the later Policy Paper, a more demanding interpretation of criteria will lead to the desired narrowing -- and perhaps eventually to a single recommended approach.