MTS Task Force Presentation

Using Risk Assessment in Inspection Programs

Douglas L. Slitor, *The Minerals Management Service*

My agency, the Minerals Management Service (MMS), has three primary responsibilities: leasing offshore lands for mineral exploration; regulating exploration, drilling, and production; and collecting royalties associated with production.

I am going to talk very directly about some ways we have been using risk assessment. Several studies in the late 1980s and early 1990s suggested that MMS should use risk assessment in some fashion in their inspection program, the program that I am in. These recommendations were made because of the following conditions. There has been offshore oil and gas exploration and production since the late 1940s, and these facilities have grown tremendously in that 50-year period. We have a wide range of facilities—nearly 4,000. Regulations require that we inspect these facilities once a year and we have a very limited inspection workforce to do that. We are barely able to get to each facility each year.

We have some very good facilities and we also have some facilities that are poor performers. We have nearly 1,500 single-well caissons, which are merely pipes sticking out of the water. We also have very large structures with 60 wells—complex equipment that can house 100 people.

Our traditional approach has been compliance strategy. We have developed 600 potential incidents of noncompliance. Our inspectors land, they go through a checklist based on the equipment there, and, if they see an infraction, they issue an incident of noncompliance (INC). We created a huge database with this information. Part of it is very useful and, as many of you are aware, part of it is faulty. But it has been useful for us.

What we want to do though, because we have an annual requirement to inspect facilities, is to focus on those facilities that are poor performers. We needed to find a way to do that. We started out by playing with semantics to try to buy ourselves some time and we redefined the term "facility." We came up with a cluster arrangement that reduced the number of inspectable facilities from almost 4,000 to fewer than 2,000.

We have yet to implement the next part, which is a sampling methodology. We have made revisions so that we can take a look at a statistically significant sampling of components and walk away with the knowledge that we have 95 percent assurance that a particular facility is in compliance. These two efforts have increased the time available to our inspectors, so hopefully we can focus on the facilities with more problems.

We started by developing a list of risk factors. We did this in concert with our inspection workforce. We did an initial survey that allowed inspectors to rank specific risk factors, which gave us a starting point for looking at our data and deciding what to do first.

After that, Dr. Paul Fischbeck introduced me to one of his graduate students, who has taken our database and taught us some things about what goes on in the database; he also developed a model that we can use to predict where accidents are likely to occur. He took 10 years worth of data (basically compliance data), INCs, infractions, accident data, and a host of other fields of information and put it into a NeuralNet software. This software learns about patterns among data and gives weights to certain things as it learns about the data. It came up with some very interesting findings.

Along with the NeuralNet, some logistical regression work was also done, and we looked at the number of INCs received by a given facility. The three correlated very strongly. The NeuralNet itself looked at data for a 5-year period, took that data, and tried to predict the facilities that were likely to have some type of incident the following year. It was fairly successful. When we rank-ordered the facilities by the risk the NeuralNet had predicted, we found that, typically, 55 percent of the accidents occurred on 20 percent of the platforms; sometimes up to 70 percent of incidents occurred on 20 percent of the platforms. This was exciting from our standpoint, because we had a method that we could use to focus our inspection workforce on some facilities that were having problems.

One of the drawbacks is that there are a lot of false positive results. Certain facilities were predicted to have accidents and they did not. What we take from this is basically that false positives might be a near-miss situation brewing. Our inspector workforce could still be validly used to check on all these facilities.

How do we plan to use these data? We have some things in mind now and we hope in 1999 to get to a point where we can run a pilot program. We want to use the NeuralNet to determine inspection frequency. It gives a value between zero and one; the higher numbers are those that the model considers more likely to have some kind of incident. We want to take that value and delineate what we believe would be a low-risk area, a medium-risk area, and a high-risk area. This tells only part of the story of the risk. A lot of things happen daily—new applications being applied for construction or welding or a particular operation where risks are higher than other operations. At that point, we want our district supervisors to take the value the model has generated and overlay it with some more specific and more recent information that they know. It may be information that is not in our database.

One of the things that the inspector workforce indicated was that we have a lot of new operators. The number of people has grown along with the number of facilities. Some are brand new and do not have a good grasp on our regulations; that is a cause for concern and it is a red flag to our workforce that is not in the database. It is something that can be used when making a final determination on what kind of inspection strategy to use for a particular platform. Also, the manpower and logistics vary—we have facilities that are very close, and we have some that are very far offshore. All this has to be factored in for the district supervisors' attention and for figuring out a particular strategy.

Where we really want to go from this point is to work with our pilot program and learn some lessons from it. We think we are on to something, but we cannot foresee the problems; we want to work this out in the field and see how it works with the inspection workforce. They have been very receptive to it because they are involved with it. They help determine the inspection frequency and the three surveys we've conducted with them indicate that they believe they have a vested interest in this.

This is a part of three programs that dovetail fairly nicely. The risk-based inspection is one. We also are working with performance measures of operators. This is facility specific, regardless of the operator, but we also look at performance of operators, and that becomes another factor that we would like to fold into this—how that operator is doing. We are beginning to talk to those operators who are doing well about alternative compliance—we are receiving proposals from them on how they can still meet the intent of the regulations but through their own means, giving them the flexibility to act on their own in terms of their own efficiencies and manpower needs.