Reliability-Centered Maintenance

By: Anthony M. Smith, PE.

Chapter 4

RCM—A Proven Approach

In this chapter, we will introduce the basic concepts that constitute what is known as *Reliability-Centered Maintenance*. Initially, however, we will briefly discuss how PM has evolved in the industrial world, and most importantly, we will look at how one of the basic tenets of reliability engineering—the "bathtub curve"—can and should influence the formulation of PM tasks. Next, we will look at how the commercial aviation industry was historically the motivating force behind the creation of the RCM methodology during the Type Certification process for the 747 aircraft in the 1960s. Finally, we will itemize the four basic features that constitute the necessary and sufficient conditions or principles that define RCM, and discuss some of the cost-benefit considerations that can accrue through the use of RCM.

4.1 Some Historical Background

If we look back to the days of the Industrial Revolution, we find that the designers of the new industrial equipment were also the builders and operators of that equipment. At the very least, they had a close relationship with the hardware that evolved from their creative genius, and as a result they truly did "know" their equipment—what worked, how well, and for how long; what broke, how to fix it, and, yes, how to take certain reasonable (not too expensive) actions to prevent it from breaking. In the beginning, then, experience did in fact play the major role in formulating PM actions. And, most importantly, these experience-based actions derived from those people who had not just maintenance experience, but also design, fabrication, and operation knowledge. Within the limits of then available technology, these engineers were usually correct in their PM decisions.

As industry and technology became more sophisticated, corporations organized for greater efficiency and productivity. This, of course, was necessary and led to numerous advantages that ultimately gave us the high-volume production capability that swept us into the twentieth century. But some disadvantages occurred also. One of these was the separation of the design, build, and operate roles into distinct organizational entities where virtually no one individual would have the luxury of personally experiencing the entire gamut of a product cycle. Thus, the derivation of PM actions from experience began to lose some of its expertise.

Not to worry! Another technology came along to help us—reliability engineering. The early roots of reliability engineering trace back to the 1940s and 1950s. Much of its origin resides in the early work with electronic populations where it was found that early failures (or infant mortalities) occurred for some period of time at a high but decreasing rate until the population would settle into a long period of constant failure rate. It was also observed that some devices (e.g., tubes) would finally reach some point in their operating life where the failure rate would again sharply increase, and aging or wearout mechanisms would start to quickly kill off the surviving population. (This scenario, of course, also very accurately describes age-reliability characteristics of the human population.) Engineers, especially in the nonelectronic world, were quick to pick up on this finding, and to use it as a basis for developing a maintenance strategy. The picture we have just described is the well-known *bathtub curve.* Its characteristic shape led the maintenance engineer to conclude that the vast majority of the PM actions should be directed to *over-hauls* where the equipment would be restored to like-new condition before it progressed too far into the wearout regime.

Thus, until the early 1960s, we saw equipment preventive maintenance based in large measure on the concept that the equipment followed the bathtub shape, and that overhaul at some point near the initiation of the increasing failure-rate region was the right thing to do.

4.2 The "Bathtub Curve" Fallacy

As this title suggests, all may not be totally well with the bathtub curve. True, some devices *may* follow its general shape, but the fact is that more has been assumed along those lines than has actually been measured and proven to be the case. As those with even a cursory knowledge of statistics and reliability theory can attest, this is not surprising, because large sample sizes are required in order to accurately develop the population age-reliability characteristics of any given device, component, or system. And such large samples, with recorded data on operating times and failures, are hard to come by.

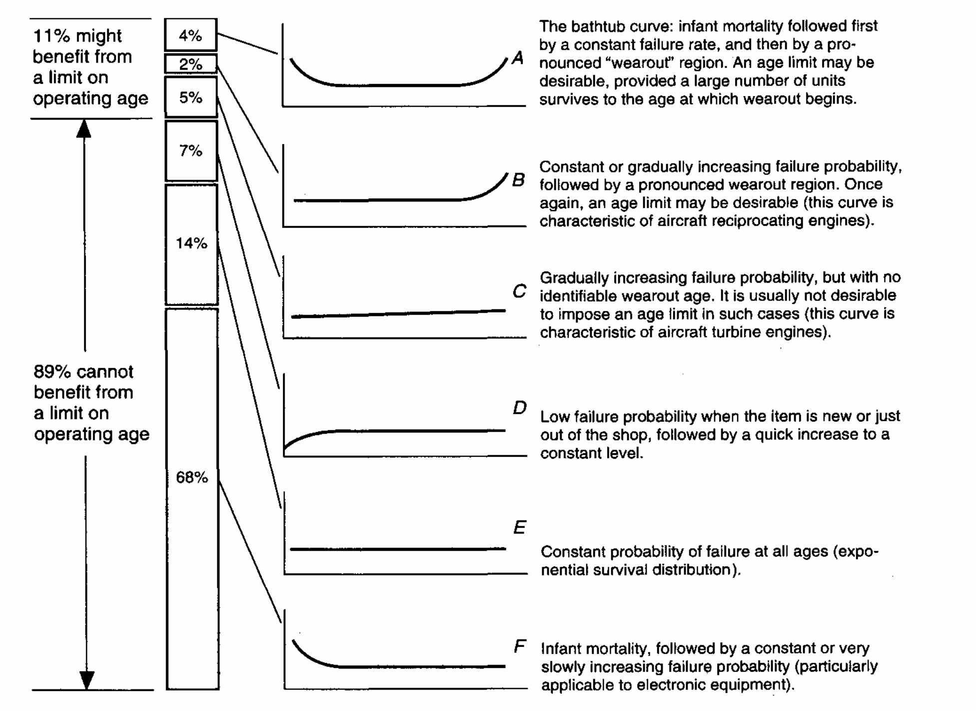

A The bathtub curve: infant mortality followed first by a constant failure rate, and then by a pronounced "wearout" region. An age limit may be desirable, provided a large number of units survives to the age at which wearout begins.

B Constant or gradually increasing failure probability, followed by a pronounced wearout region. Once again, an age limit may be desirable (this curve is characteristic of aircraft reciprocating engines).

C Gradually increasing failure probability, but with no identifiable wearout age. It is usually not desirable to impose an age limit in such cases (this curve is characteristic of aircraft turbine engines).

D Low failure probability when the item is new or just out of the shop, followed by a quick increase to a constant level.

E Constant probability of failure at all ages (exponential survival distribution).

F Infant mortality, followed by a constant or very slowly increasing failure probability (particularly applicable to electronic equipment).

Figure 4.1 Age–reliability patterns for nonstructural aircraft equipment (United Airlines).

The commercial aviation industry, however, does have fairly large populations of identical or similar components in their aircraft fleets— components that are common to several aircraft types. And, as an industry, they have made some deliberate and successful efforts to accumulate a database of operating history on these components. Such a database is driven by several factors, not the least of which are safety and logistics considerations. As a part of the extensive investigation that was conducted in the late 1960s as a prelude to the RCM methodology, United Airlines used this database to develop the age-reliability patterns for the nonstructural components in their fleet. This was done as a part of the more general questioning that preceded RCM concerning whether airline equipments did, in fact, follow the bathtub curve. Specifically, failure density distributions were developed from the component operating history files, and the hazard rate (or instantaneous failure rate) was derived as a function of time. The results of this analysis are summarized in Fig. 4.1.

These results came as a surprise to almost everyone—and continue to do so even today when people see these results for the first time. The significance of these results, and their potential importance to the maintenance engineer, cannot be stated too strongly. Let's examine these more closely, assuming for the moment that these curves may be characteristic of your plant or system:

1. Only a very small fraction of the components (4 percent) actually replicated the traditional bathtub curve concept (curve *A).*
2. More significantly, only 6 percent of the components experienced a distinct aging region during the useful life of the aircraft fleet (curves *A* and *B).* If we are generous in our interpretation, and allow that curve *C* also is an aging pattern, this still means that only 11 percent of the components experienced an aging characteristic!
3. Conversely, 89 percent of the components never saw any aging or wearout mechanisms developing over the useful life of the airplanes (curves *D, E,* and *F).* Thus, while common perceptions tend toward the belief that 9 of 10 components have "bathtub" behavior, the analysis indicated that this was completely reversed when the facts were known.
4. Notice that 72 percent of the components, however, did experience the infant mortality phenomenon (curves *A* and *F).*
5. And the most common grouping, 68 percent, was starting to look like the bathtub, but never got to the aging region (curve *F).*

What does all of this mean? Quite a bit! First, recall that a constant failure rate region (curves *A, B, D, E,* and *F* all have this region) means that the equipment failures in this region are random in nature—that is, the state of the art is not developed to the point where we can predict what failure mechanisms may be involved, nor do we know precisely when they will occur. We only know that, on average in a large population, the hazard rate (or the mean time between failure) is a constant value. Of course, we hope that this constant-failure-rate value is very small, and we thus have a very reliable set of components in our system. But, for the maintenance engineer, these constant-failure-rate regions mean that overhaul actions will essentially (short of luck) do very little, if anything, to restore the equipment to a like-new condition. In this constant-value region, overhaul is usually a waste of money because we really do not know what to restore, nor do we really know the proper time to initiate an overhaul. (In the constant-failure-rate region, any time you might select is essentially the wrong time!) Second, and worse yet, is that these overhaul actions will actually be harmful because, in our haste to restore the equipment to new, pristine conditions, we have inadvertently pushed it back into the infant-mortality region of the curve. In this specific study, for example, over-haul actions on 72 percent of the components (curves *A* and *F)* would be susceptible to this counterproductive situation. A third point relates to the periodicity that should be specified for an overhaul task when such an action is considered to be the correct step to take. For example, if a component is either a curve *A* or *B* type, we want to assure that the overhaul action is not taken too soon—or again, we may be wasting our resources.

Often, we do not know what the correct interval should be, or even if an overhaul PM task is the right thing to do. Why? Because we do not have sufficient data to tie down the age-reliability patterns for our equipment. In these instances, we may wish to initiate an Age Exploration program.

In summary, we should be very careful about selecting overhaul PM tasks because our equipment may not have an age-reliability pattern that justifies such tasks. In addition, overhauls are likely to cause more problems than they prevent if aging regions are not present. When data is absent to guide us on this very fundamental and important issue, we should initiate an Age Exploration program and/or the collection of data for statistical analyses that will permit us to make the right decisions. It is indeed a curious (and unfortunate) fact that in today’s world of modern technology, one of the least understood phenomenon about our marvelous machines is how and why they fail!

4.3 The Birth of RCM

RCM epitomizes the old adage that “necessity is the mother of invention.”

In the late 1960s, we found ourselves on the threshold of the jumbo jet aircraft era. The 747 was no longer a dream; the reality was taking shape as hardware at the Boeing factory in Seattle. The licensing of an aircraft type (called *Type Certification* by the FAA) requires, among its many elements, that an FAA-approved preventive maintenance pro-gram be specified for use by all owners/operators of the aircraft. No air-craft can be sold without this Type Certification by the FAA. The recognized size of the 747 (three times as many passengers as the 707 or DC-8), its new engines (the large, high bypass ratio fan jet), and its many technology advances in structures, avionics, and the like, all led the FAA to initially take the position that preventive maintenance on the 747 would be very extensive—so extensive, in fact, that the airlines could not likely operate this airplane in a profitable fashion. This development led the commercial aircraft industry to essentially undertake a complete reevaluation of preventive maintenance strategy. This effort was led by United Airlines who, throughout the 1960s, had spear-headed a complete review of why maintenance was done, and how it should best be accomplished. Names like Bill Mentzer, Tom Matteson, Stan Nowland, and Harold Heap, all of United Airlines, stand out as the pioneers of this effort. What resulted from this effort was not only the thinking derived from the curves in Fig. 4.1, but also a whole new approach that employed a decision-tree process for ranking PM tasks that were necessary to preserve critical aircraft functions during flight. This new technique for structuring PM programs was defined in MSG-1 (Maintenance Steering Group-1) for the 747, and was subsequently approved by the FAA. The MSG-1 was able to sort out the wheat from the chaff in a very rational and logical manner. When this was done, it was clear that the economics of preventive maintenance on a 747-sized aircraft were quite viable—and the 747 became a reality.

The MSG-1 was so successful that its principles were applied in MSG-2 to the Type Certification of the DC-10 and L-1011. In recent times, MSG-3 has developed the PM program for the 757 and 767. Versions of the MSG format have likewise guided the PM programs for the Concorde, Airbus, 737-300/400/500, and various retrofits to aircraft such as the 727-200, DC-8 stretch, and DC-9 series.

In 1972, these ideas were first applied by United Airlines under Department of Defense (DOD) contract to the Navy P-3 and S-3 air-craft and, in 1974, to the Air Force F-4J. In 1975, DOD directed that the MSG concept be labeled “Reliability-Centered Maintenance,” and that it be applied to all major military systems. In i978, United Air-lines produced the initial RCM “bible” under DOD contract.

Since then, all military services have employed RCM on their major weapons systems. RCM specifications have been developed, the Air Force Institute of Technology (AFIT) offers a course in RCM, and the Navy has published an RCM handbook.

The most recent use of RCM has occurred in the utility (electric power generation) industry. In 1983, the Electric Power Research Institute (EPRI) initiated RCM pilot studies on nuclear power plants. Since these early pilot studies, several full-scale RCM applications have been initiated in commercial nuclear and fossil power plants.

Clearly, the development of RCM has been an evolutionary process, and some 30 years have passed during which RCM has become a mature process in commercial aviation and DOD, and an embryo process in power generation plants. Basically, no other industry has yet fully embraced RCM, in spite of its proven track record. Hopefully, this book will help to change that picture.

4.4 What Is RCM?

Some of the more prominent practices and myths that currently constitute the basis for PM program development can be summarized by saying that these practices and myths are driven, in large measure, by the overriding consideration and principle of “what can be done?”—but rarely by traceable decisions such as “why should it be done?” Stated another way, we could say that the overriding motivation of current PM practices is to “preserve equipment operation.” Until recently, this has resulted in little, if any, consideration as to why we take certain PM actions and what, if any, priority should be assigned to the expenditure of PM resources. Almost without fail, maintenance planning starts directly with the equipments and seeks to specify as quickly as possible the various things that are felt necessary to “keep it running” (sometimes without a conscious evaluation of the function that is served or consideration of a cost-benefit decision).

On the other hand, RCM is not just another way to do this same old thing all over again. It is very different in some fundamental aspects from what is today the norm among maintenance practitioners, and requires that some very basic changes in our mind-set must occur. As you will see in a moment, the basic RCM concept is really quite simple, and might be characterized as organized engineering common sense.

There are four features that define and characterize RCM, and set it apart from any other PM planning process in use today. Each of these is defined and discussed here.

Feature 1 . The most important of the four RCM features is also per-haps the most difficult to accept because it is, at first glance, contrary to our ingrained notion that PM is performed to preserve equipment operation. *The primary objective of RCM is to preserve system function.* Notice that this objective does not initially deal with preservation of equipment operation. Of course, we will ultimately preserve system function by preserving equipment operation, but not as the initial step in the RCM process. By first addressing system function, we are saying that we want to know what the expected output is supposed to be, and that preserving that output (function) is our primary task at hand. This first feature enables us to systematically decide in later stages of the process just what equipments relate to what functions, and will not assume a priori that "every item of equipment is equally important," a tendency that seems to pervade the current PM planning approach.

Let's look at a couple of simple examples to illustrate the inherent value associated with the "preserve system function" concept. First, compare two separate fluid transfer trains in a process plant where each train has redundant legs. Train A has 100 percent capacity pumps in each leg, and train B has 50 percent capacity pumps in each leg. As the plant manager, I tell you, the maintenance director, that your budget will allow PM tasks on either train A pumps or train B pumps, but not both. What do you do? Clearly, if you don't think function, you are in a dilemma, since your background says that your job is to keep all four pumps up and running. But if you do think function, it is clear that you must devote the defined resources to the train B pumps since loss of a single pump reduces capacity by 50 percent. Conversely, a loss of one pump in train A does not reduce capacity at all, and also in all likelihood allows a sizable grace period to bring the failed pump back to operation. As a second example, let's examine more closely just what function is really performed by a pump. The standard answer is to preserve pressure or maintain flow rate—a correct answer. But there is another, more subtle, function to preserve fluid boundary integrity (a passive function). In some cases, allocation of limited resources to PM tasks for the passive function could be more important than keeping the pump running (e.g., when the fluid is toxic or radioactive). Again, if you don't think function, you may miss drawing the proper attention to the passive boundary integrity function.

Feature 2. Since the primary objective is to preserve system function, then *loss of function or functional failure* is the next item of consideration. Functional failures come in many sizes and shapes, and are not always a simple "we have it or we don't" situation. We must always carefully examine the many in-between states that could exist, because certain of these states may ultimately be very important. The loss of fluid boundary integrity is one example of a functional failure that can illustrate this point. A system loss of fluid can be (1) a very minor leak that may be qualitatively defined as a drip, (2) a fluid loss that can be defined as a design basis leak—that is, any loss beyond a certain GPM value will produce a negative effect on system function (but not necessarily a total loss), and (3) a total loss of boundary integrity, which can be defined as a catastrophic loss of fluid and loss of function. Thus, in our example, a single function (preserve fluid boundary integrity) led to three distinct functional failures.

The key point in feature 2 is that we now make the transition to the equipment by identifying the *specific failure modes* in specific components that can potentially produce those unwanted functional failures.

Feature 3. In the RCM process, where our primary objective is to pr serve system function, we have the opportunity to decide, in a very systematic way, just what order or priority we wish to assign in allocating budgets and resources. In other words, "all functions are not created equal," and therefore all functional failures and their related components and failure modes are not created equal. Thus, we want to *prioritize the importance of the failure modes.* This is done by passing each failure mode through a simple, three-tier decision tree which will place each failure mode in one of four categories that can then be used to develop a priority assignment rationale.

Feature 4. Notice that up to this point, we have not yet dealt directly with the issue of a preventive maintenance action. What we have been doing is formulating a very systematic roadmap that tells us the where, why, and priority with which we should now proceed in order to establish specific PM tasks. We thus address each failure mode, in its prioritized order, to identify candidate PM actions that could be considered. And here, RCM again has one last unique feature that must be satisfied. Each potential PM task must be judged as being "applicable and effective." *Applicable* means that if the task is performed, it will in fact accomplish one of the three reasons for doing PM (i.e., prevent or mitigate failure, detect onset of a failure, or discover a hidden failure). *Effective* means that we are willing to spend the resources to do it. Generally, if more than one candidate task is judged to be applicable, we would opt to select the least expensive (i.e., most effective) task. Recalling, that when describing a run-to-failure task category, there are three reasons for making such a selection - failure of a task to pass either the applicability or effectiveness test results in two of the run-to-failure decisions, while the third would be associated with a low-priority ranking and a decision not to spend PM resources on such insignificant failure modes.

In summary, then, the RCM methodology is completely described in four unique features:

1. Preserve functions.
2. Identify failure modes that can defeat the functions.
3. Prioritize function need (via the failure modes).
4. Select only applicable and effective PM tasks.

4.5 Some Cost-Benefit Considerations

As noted earlier, the primary driving force behind the invention of RCM was the need to develop a PM strategy that could adequately address system availability and safety without creating a totally impractical cost requirement. This has clearly been successfully achieved by commercial aircraft; however, quantitative data in the public arena on this cost picture is rather scarce.

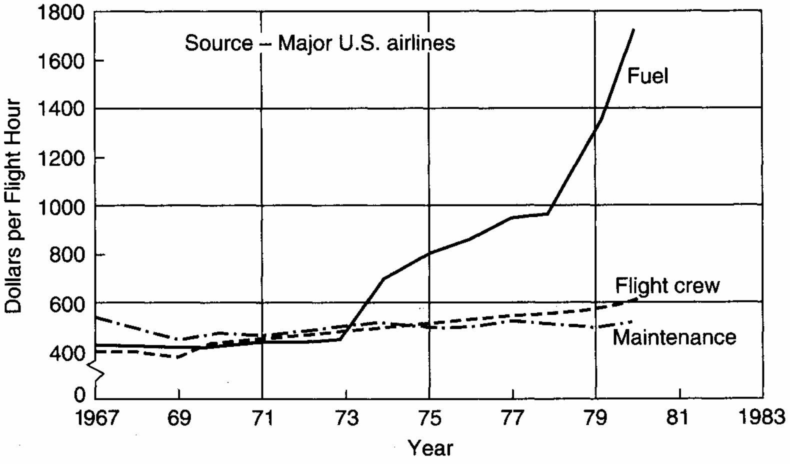

Figure 4.2 Costs per flight-hour (1982 constant dollars).

Figure 4.2 presents maintenance cost per flight hour in the first 10 years of RCM use. What we see in Fig. 4.2 is a maintenance cost that is essentially constant from the late 1960s to the early 1980s. This is precisely the period during which the 747, DC-10, and L-1O11 were introduced into revenue service on a large scale. These jumbo jets not only introduced the large increase in passenger capability per aircraft (about 3 times over the 707 and DC-8), but also a higher daily usage rate and the deployment of several advanced technologies into everyday use. In spite of all of these factors, any one of which would normally tend to drive maintenance costs up, we see a fairly constant maintenance cost per flight hour historically occurring. RCM was the overriding reason for this.

Figure 4.3 presents another way to view the impact of RCM in the commercial aircraft world. Note that the PM definitions used in Fig. 4.3 correspond as follows to our set of PM task definitions:

Hard-time	Time-directed
On-condition	Condition-directed
Condition-monitored	Run-to-failure

Two significant points can be observed with this data. First the pre- (1964) and post- (1969/1987) RCM periods reveal the dramatic shift that occurred in the reduction of costly component overhauls (i.e., hard-time tasks), mainly in favor of run-to-failure (i.e., condition-monitored) tasks. Much of this shift, of course, was made possible by a design philosophy that included double and triple redundancy in the flight-critical functions. The RCM process was employed to take advantage of these design features in structuring where PM was critical and where the run-to-failure decision was appropriate. Also, notice that throughout the time period represented here, the condition-directed (i.e., on-condition) task structure remained fairly constant. The commercial aircraft industry was one of the early users of performance and diagnostic monitoring as a PM tool, and they have continued to successfully apply it throughout the generation of the newer jet aircraft.

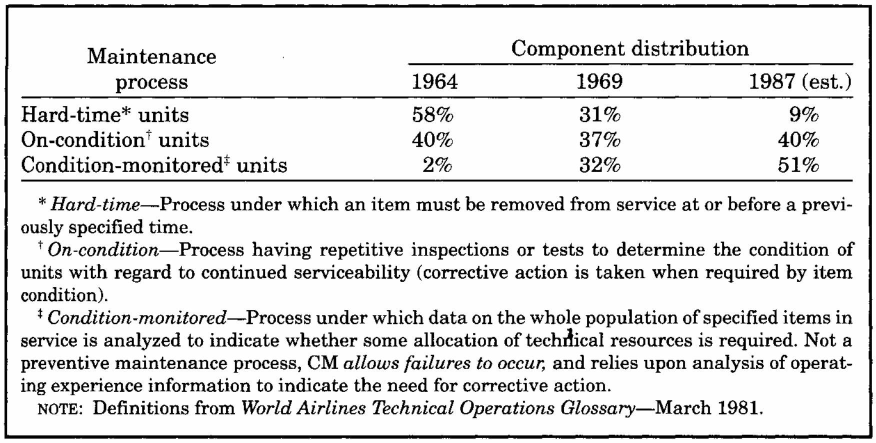

Maintenance process	Component distribution		
	1964	1969	1987 (est.)
Hard-time* units	58%	31%	9%
On-condition† units	40%	37%	40%
Condition-monitored‡ units	2%	32%	51%

* *Hard-time*—Process under which an item must be removed from service at or before a previously specified time.

† *On-condition*—Process having repetitive inspections or tests to determine the condition of units with regard to continued serviceability (corrective action is taken when required by item condition).

‡ *Condition-monitored*—Process under which data on the whole population of specified items in service is analyzed to indicate whether some allocation of technical resources is required. Not a preventive maintenance process, CM *allows failures to occur,* and relies upon analysis of operating experience information to indicate the need for corrective action.

NOTE: Definitions from *World Airlines Technical Operations Glossary*—March 1981.

Figure 4.3 Commercial aircraft—component maintenance policy.

The results indicated in Figs. 4.2 and 4.3 have led to a growing interest in other commercial areas. Most notably, nuclear power generation plants in several U.S. utilities are currently implementing RCM—or, at the very least, are conducting RCM pilot projects as a prelude to the conduct of a comprehensive RCM effort. Foreign nuclear utilities are also pursuing RCM, and Electricité de France (EDF) has in fact embarked upon a major effort to incorporate RCM as the basis for their PM program in all 50 operating nuclear units. Fossil power generation units in the United States have also started to apply RCM. In Chap. 7, selected information from RCM programs at GPU Nuclear Corporation, Florida Power & Light Company, and Westinghouse Electric Corporation will be presented to illustrate some typical RCM results. The reasons for this interest among the utilities and power plant operators are cost driven from two points of view:

1. Control and reduction of O&M costs
2. Increase in plant availability

The first RCM pilot study in 1984 at FP&L's Turkey Point Nuclear Plant provided the initial evidence of how an RCM program could favorably impact the O&M cost picture. Analysis showed that a 30 to 40 percent savings in the existing PM program cost for the component cooling water system could be realized by implementing the RCM-based results. More recent surveys conducted by the Electric Power Research Institute have continued to quantify and verify this initial conclusion. In the EPRI survey involving seven different utilities, it was found that an average cost payback period of 6.6 years had been demonstrated with early RCM programs. This data was then extrapolated to a mature program state to show that the cost payback period could readily be reduced to 2 years or less—and this on the basis of PM cost savings alone. Other O&M benefits yet to be evaluated and credited to the cost-benefit picture include the following:

- plant trip reductions
- documented basis for the PM programs for use with regulators, insurers, and warranty contracts
- more accurate spare parts identification and stocking
- more efficient PM planning and scheduling
- decrease in corrective maintenance costs

But perhaps the most dramatic cost savings will be realized in the plant availability area where cost-avoidance benefits will be very large. For example, a base-loaded generation unit in the 1000 MWe range costs about $500,000—$750,000 for replacement electricity for *each* day of a forced outage. Thus, an RCM-based PM program that could raise plant availability by only one or two percentage points has a payoff in the multimillion-dollar range.

All of this is to say that the cost-benefit payoff with RCM has been dramatic in its impact on commercial aviation, and potentially offers similar dramatic payoffs in other areas where complex plants and systems are routinely operated.