How To Succeed At PM Improvement

David H. Worledge, Asset Performance Technologies, Inc.

3rd Annual SMRP Houston Chapter
Maintenance & Reliability Symposium
August 27th, 2009

1. Background in Reliability Centered Maintenance

For about a decade or more from the early 80's through the mid 90's the nuclear power industry undertook the daunting task of rationalizing PM programs that across the industry had grown very large and costly but yet still lacked a recognizable technical basis. I was privileged to play a central role in that effort by managing the drive to implement Reliability Centered Maintenance (RCM), sponsored by the Electric Power Research Institute. Ultimately, that effort failed to demonstrate that RCM is a practical technique for large land-based industrial facilities, but it did succeed resoundingly in learning how to craft a high quality approach that is reliability-based, and sufficiently efficient for companies to sustain over the long haul. This paper is a summary of those lessons learned and the methodology that resulted.

Oil and gas installations share with nuclear power plants defining characteristics that are important to maintenance. They both comprise, 1) a large number (several tens of thousands) of industrial infrastructure equipment items (e.g. pumps, motors, turbines, instruments, valves, heat exchangers, switchgear) with a much smaller number of major equipment items (gas turbines, fractionating columns) specific to the production process, 2) have large downtime costs, 3) manipulate hazardous materials with high penalties for safety and environmental infringements, and 4) have high capacity factors with long periods of continuous production between major turnarounds. Of course, there are also differences, but both logic and our experiences dictate that the problems of PM improvement in large installations are ubiquitous and the same solution found in nuclear power will be successful in the oil and gas industry.

2. Objectives

The aim of this paper is to demonstrate the essential requirements for efficient completion of PM improvement efforts, while retaining the reliability basis introduced originally by RCM. A reliability basis means that each PM activity has to be focused very clearly, and in significant detail, on the known ways in which the equipment degrades, and the timing of the PM activities has to be logically related to the patterns and time scales of the degradation. Economic efficiency in program development and subsequent program implementation imposes strong constraints on when PM is needed and when it is not. The objective of this paper is to bring these major driving forces to your attention and to describe how they can best be managed without sacrificing the quality of the results. These forces can be summarized as the conflicts between the need for technical detail and rigor and the logistical demands of providing it for the bulk of the equipment in scope.

3. The Challenge

3.1 Scale

I interpret PM to include, 1) scheduled tasks such as inspections, calibrations, and refurbishments, 2) predictive tasks such as vibration analysis, various non-destructive testing methods, and oil sampling and analysis, 3) lubrication, 4) failure-finding activities such as functional tests, 5) end-of life replacements, and 6) frequent tasks usually carried out by non-maintenance personnel such as operator rounds and system or machinery owner walkdown inspections. We might end up with an average of 3 to 5 PM tasks per equipment item, sometimes a lot more.

To establish a baseline for resource estimation let us assume that the most significant PM program issues at a typical plant can be addressed by analysis of 10,000 equipment items. The tables of degraded conditions (loosely referred to as FMEAs) will average a few dozen rows for each item. It is obvious that we can easily reach over a million (3x10,000x50) mappings of PM tasks to degraded conditions, and the FMEA tables and suitable PM tasks must be developed in the first place. If each degraded condition and PM mapping consumes even a fraction of a minute of time we will need about 10,000 hours just for this part of the job – and we have not yet multiplied by the several people normally needed to do this analysis, most of them from your own plant. Clearly we must group equipment items together on some basis, develop FMEAs that we use over and over, and be very systematic about the way we apply PMs to equipment IDs.

3.2 Implementation

Regardless of how the PM evaluation is carried out, the results must be translated into data in the plant's enterprise data system. It turns out that this can be a more daunting task than the PM evaluation, however the latter is done. Implementation can cut across organization boundaries thus requiring diplomacy and teamwork, sometimes not in step with the pervading culture. It is also information and format intensive. Although we do not have time to examine this aspect in detail in this paper the following guidelines can help you be successful. Ignoring any one of these items sets you up for failure with a high probability.

1. You need a detailed implementation plan right from the start of the PM evaluation effort. The job is too production oriented to be finding out essential requirements as you go. The plan should address all of the items below, including exactly which personnel and which organizations will be involved.

2. Find out your company's corporate PM implementation requirements as soon as possible. These easily might not be known at the plant. You may have to dig for this information; better sooner than later. Be prepared for regulations on risk estimation and document formats, and IT restrictions on access to the enterprise data system, especially for contractors.

3. Get buy-in *from all parties affected* on the degree of technical detail required in implementation documents such as task sheets, job plans, and formal and informal procedures. Nail down the details,

including the way these documents are to be formatted, with an agreed set of example documents before you start to implement anything. Since this can take a long time in contentious cases, get started as early as possible.

4. Figure out the most effective review and sign-off process, including the considerable demand on plant personnel even when the bulk of implementation work is to be done by contractors. For a large production job like this you may need to negotiate a different process than is customary for PM changes done on a day to day basis.

5. Implement PM evaluation results as soon as they are approved. Do not wait until the evaluation phase is completed.

3.3 Costs, Benefits, And Expectations

It seems obvious that senior management should not expect to operate industrial plant cost-effectively without a well developed PM program. However, this does happen more often than you would think. Even when the PM program functions reasonably well, senior management may not be familiar with the benefit-to-cost ratio of improving it. A moment's reflection shows that with about fifty to a hundred ways a piece of complex machinery will degrade in service, its MTBF will have a strong dependence on whether you take steps to find severely degraded conditions before they find you. It is quite common for a good PM program to reduce the failure rate of such equipment by a factor of ten or more. The average annual cost of such a PM program per equipment ID is a few hundred to a few thousand dollars, whereas an in-service failure can easily cost over a million dollars in lost production or other penalties. Effective PM therefore has a high benefit to cost ratio when *total* costs are considered, typically exceeding a factor of ten and sometimes a lot more. Where else can a plant manager get that rate of return for the company? However, given the logistics, the secret is to not go broke trying to achieve a good PM program. Of course, going broke is not an option, but giving up on the improvement effort is a common experience, and results in foregoing the much higher productivity and safety that a better program would achieve. Any large scale PM improvement must be expected to cost a lot in time and effort, so the focus must be on efficiency and value-added, as well as technical quality.

3.4 Overview of Recommended Approach

Traditional RCM was performed by a team of plant personnel with each equipment item studied individually, each in its own operating context. This works well when the analysis can be performed once with the costs reduced by being spread over a large number of identical plants (e.g. a fleet of identical airplanes) and the benefits multiplied by the same number. It is not generally effective for unique plant designs except in special problem-solving circumstances. Operating context means the combinations of functional importance, duty cycle, and service conditions that apply to an individual instance (i.e. equipment ID) of an equipment type (e.g. multistage horizontal pump). The methodology described in more detail in the following sections is illustrated in Figure 1 and consists of:

1. Develop only one FMEA for each equipment type that covers the full range of operating contexts.

2. PM strategies (e.g. Vibration Analysis, Calibration, Overhaul) are not arrived at by being assembled bottom-up from individual actions prescribed to address the degraded conditions. Instead they are

assigned at the start because they are well known for standard types of equipment. Individual line items within each PM strategy are subsequently identified by top-down mapping within the FMEA to individual degraded conditions.

3. Equipment ID's are grouped into PM Groups as to type and operating context in a standard format.

4. The PM strategies from the FMEA, now including task intervals, are organized into a Template using the same standard format of the operating context.

5. PM tasks and intervals are assigned from the Template to one equipment member of each PM group of equipment ID's according to the operating context. The assignments are then copied across all equipment ID's in the PM Group.

4. Essential Elements

4.1 Equipment List

A list of in-scope equipment is an obvious requirement for this analysis. Spreadsheets or database tables are the usual formats for handling steps in the analysis, all of which use the same list. Preparation of the items in the equipment list to be complete, accurate, and unique can waste an extraordinary amount of time right at the start of the analysis. It does not take the expertise of PM analysts to do this, so plant personnel can reduce wasted time and expense by preparing the equipment list in advance of work by contractors. Duplicate, missing, and erroneous equipment ID's, equipment abandoned in place, and equipment "parented" by other equipment must be addressed.

4.2 Criticality

Conceptually, cost-benefit considerations require a higher level of PM resources to be devoted to equipment that costs a lot when it fails in-service than to equipment that has a less costly impact. In principle this should be addressed by considering both the frequency and consequences of failure. Both RCM and the approach adopted by the nuclear power industry specifically address only consequences. Frequency is dealt with implicitly by including in the FMEA only those degraded conditions known to occur, acknowledging that most equipment types have MTBF's in the range of 1-10 years (much narrower than the range of consequences), and are present in sufficient numbers and for sufficient time to ensure that if they can fail they will. This is a very valuable and valid simplification.

Although management may have reasons to prefer consequences of failure to be described by a dozen or more categories (and it is not too time consuming to do this at the equipment level) it is not strictly important to describe consequences in any more detail than is required to identify truly major differences in the level of PM to be applied. However many consequence categories are identified, they feed into three levels of functional importance that are commonly used, 1) Run-To-Failure (RTF), which means apply no PM at all, 2) Non-Critical (i.e. of minor functional importance), and 3) Critical (i.e. where the cost of in-service failure is very high). It is common to find the Critical level split into two, Critical A and Critical B. This distinction is most useful for separating critical equipment protected by equipment-level redundancy (Critical B), from that which is not (Critical A). Equipment redundancy is worthless unless the backup item works when required, which is not usually well assured if a Non-Critical level of

PM is applied. However, it is clearly not necessary to aim for redundant equipment to be as reliable as non-redundant equipment, even if the function provided is highly critical. A reasonable way to think of the objectives of PM is that you are trying to prevent all failures of critical A equipment, most failures of critical B equipment, and some failures of non-critical equipment. I prefer to use just Critical, Non-Critical, and RTF because Critical/Non-Critical captures the major effects using a binary criticality choice that is symmetric with the treatment given to duty cycle and service conditions, thus preserving a simple structure for the PM Template (see below). RTF, of course, does not impact the Template at all.

The process of identifying the relevant consequences of in-service failure can be carried out by a combination of an operator and a process engineer working through the equipment list with the help of a facilitator. Their knowledge is extensive and sufficient. This process is rapid and can be completed early in the project. Assignment of the criticality in this way sets priorities among many other tasks. In particular resist all attempts to force you into a complex top-down functional analysis of the kind required by RCM. This is too resource intensive, takes too long, and does not repay the extra effort involved. The RCM functional analysis was one of the first parts of RCM to be abandoned by the nuclear industry.

4.3 Equipment Types and Operating Context

The methodology recommended here depends on developing detailed FMEAs for more or less generic equipment types (e.g. multistage horizontal pump) in a way that enables them to be used in different operating contexts for individual instances of such types. An oil refinery may have 25 similar charging pumps, some operating all the time, some in standby, with individual pumps normally experiencing a variety of service stressors (high speed, high vibration, corrosive fluids, etc). *You want to do the FMEA only once*, noting the internal variabilities introduced by a *standard set* of operating contexts. It turns out that this is surprisingly easy to do when approached in the right manner. You have to spend some time early in the evaluation classifying your equipment in terms of the standard operating contexts, but you will need to do this anyway as far as criticality is concerned. Other specifications such as design type, lubrication type, speed, and service stressors would have to be addressed anyway, even if you had chosen to tackle the FMEA individually for each pump ID.

4.4 Degradation Mechanisms

These follow standard practice in FMEA construction. However a few details are worth noting. FMEA is known as failure mode and effects analysis. In the context of design reliability and safety analysis work failure *mode* has usually meant fails to start, fails to run, fails to open etc. This is not at all adequate for maintenance analysis because such modes can be contributed to by a wide variety of degraded conditions that must be individually addressed. Many such conditions can contribute to more than one failure mode (e.g. inadequate valve stem lubrication can cause a valve to fail open or closed). Maintenance analysis has to address specific degraded conditions, and these must be described at a level of cause sufficient to identify the failure pattern and time scale of failure. A good scheme is to identify the hardware involved (the subcomponent that is the site of degradation), the mechanism of degradation (e.g. wear, fatigue, change of material properties), and the driving influence or cause (e.g. normal use, high vibration, heat). The latter step is essential because without it the pattern of failure (e.g. wearout), and the time scale (e.g. expect the earliest failures in 2 years) cannot be discerned.

It is a good idea to adopt a simple coding scheme for failure pattern (e.g. W for wearout, R for random), and for the earliest time to failure for mild service conditions and low duty cycle. Adopt a similarly simple scheme for the effect of more strenuous duty cycle and service conditions and the factor by which these could affect the earliest time to failure.

Adding this information will take a little longer than a standard FMEA, but notice that we have saved a lot of time by not including a description of consequences at the row level in the FMEA. It is important to keep consequence assignment at the equipment level as described in 4.2. This is discussed further under Risk, below.

4.5 PM Mapping To Degradation Mechanisms

You need to bring a similar level of organization to the assignment of PM tasks or strategies. Continuing with the pump example, there are only a few PM strategies suitable for pumps of this kind – and they are well known, at least by those skilled in the art. Introduce each one as a separate column in the equipment type FMEA. Use the column to map the PM strategy to the degraded conditions it addresses. This mapping is independent of operating context, *so you do it only once*. The process automatically generates the line items that constitute a summary of the actions required when the task is performed (i.e. inspect the…. for….). Practitioners of RCM refer to assembling these line item instructions into PM strategies as "task packaging". Identification of the major PM strategies at the start thus *avoids the need for separate task packaging*.

4.6 PM Templates

The PM Template is just a table with the PM strategies in the leftmost column and corresponding task intervals in the remaining columns according to the format chosen for the operating context. It has been stated that a convenient format is to adopt a binary choice for each of: criticality (critical and non-critical), duty cycle (high and low), and service conditions (severe and mild), requiring 8 ($=2^3$) columns for the intervals. The Template then can be used for an operating context corresponding to any combination of these three quantities (e.g. critical, high duty cycle, and severe service conditions, through to the other end of the range, i.e. non-critical, low duty cycle, and mild service conditions).

We prefer to complete recommendations for the task intervals in the Template using the panel of plant subject matter experts and later cross check with the degradation time scales that each PM strategy is supposed to address, as mapped in the FMEA. This can be illuminating. A given Template interval may not correspond exactly to the shortest time to failure supposedly addressed by that PM strategy in the FMEA. This poses a dilemma. Selecting an interval to address the shortest time to failure could appear to be conservative when most of the other FMEA rows have much longer time scales. However, short time scales clearly imply higher failure rates if they are not addressed, so they should be respected for the critical columns in the Template. Such cases might nevertheless be assigned somewhat less stringent task intervals for Critical B compared to Critical A, if you bother to treat these cases separately, and the intervals can certainly be relaxed more for the non-critical cases. If you go further to show uncertainty in the wearout times in the FMEA, these too can be used to assess limits on task intervals. Adopting these or similar practices provides some satisfying consistency in the process of deciding appropriate task intervals across the range of operating contexts in the presence of uncertainty. For example, compare this process with trying to decide intervals for PM actions to address each row in the FMEA, one at a time, especially if you developed a new FMEA for each equipment item. It is almost impossible to maintain consistent treatment of PM task intervals in those circumstances.

4.7 Treatment of Risk

The JA1011 RCM standard discusses consequence assignment in a way that does not seem to enforce a treatment of consequences at the level of each row of the FMEA. This appears to acknowledge the time consuming nature of the activity. Operators of nuclear power plants address the consequences of failure only at the level of the equipment ID as a whole. Unquestionably, there are cases where the importance of a failure mode (in the safety analysis meaning of the term, e.g. external leak) could differ markedly in consequences from, for example, failure to run. However, consider these facts: 1) the failure *mechanisms* of the FMEA suitable for maintenance analysis are at a more detailed level than these failure *modes*. So even if addressing consequences at the mode level were important, this still is a long way from enforcing treatment for each row of the FMEA; 2) the failure modes are few in number and most of them are irrelevant from the point of view of maintenance as noted above (and simply mean the equipment will not work as intended); and 3) the only failure modes that stand out are leakage and a high energy event such as sudden rupture. These will both be easily recognized by the analysts at the level of the equipment as a whole.

Assigning consequences and thus criticality to the whole equipment as recommended here, means that the equipment could be classed as critical on the basis of leakage or an energetic destructive event that arises from only a small minority of degraded states, and thus be driven into a comprehensive PM program when perhaps such conditions could be controlled by a more focused and less expensive PM program. First, these cases are not common, and the economy in analysis time for the whole plant typically far exceeds the cost of the extra PM tasks in a few rare cases. Second, it is not advisable in such cases to rely too much on treating individual FMEA rows as independent of the others. High energy equipment experiences strong internal coupling between many degradation mechanisms due to complex thermal, hydrodynamic, and mechanical effects. If a minority of degraded conditions can lead to energetic consequences, it is not too conservative to treat the whole equipment as if it is critical.

In addition to saving analysis effort, keeping the consequence treatment at the equipment level instead of taking it down into the FMEA means the FMEA can remain application invariant, and thus enables you to develop only one general purpose FMEA for each equipment type. The application variables come with the individual instances to which it is applied. It also means that even if you do feel impelled to quantify risk in terms of both frequency and consequences, this need be done only once for each equipment ID. The methodology does not prevent you from using a Risk Matrix, but it is a peripheral issue and not essential from a maintenance perspective.

Some companies require a PM change to be justified by showing a significant decrease in risk between the unmitigated case and the mitigated case. For several reasons this needs to be handled at the equipment level by showing that the whole PM program for the equipment provides a risk reduction, and even then the frequency bins employed in a risk matrix intended for higher level applications may be too coarse for PM justification. It will not work well for individual PM tasks because, 1) they do not often provide sufficient frequency reduction by themselves, and 2) the degree of frequency reduction depends on what other PM tasks are being done, so the result of the mitigation estimate depends on factors other than the task itself.

5. Conclusions

There are many pitfalls to performing PM improvement efficiently enough to be able to sustain the investment in time and effort long enough to get the job done. At every turn it requires the utmost care be put into the methodology to speed the analysis while retaining technical quality. From this paper you will have learned about several parts of the analysis that can be greatly speeded up, and that several of these are greatly influenced by the level of granularity at which they are addressed. You should also have learned that when you take short cuts they need to be very carefully justified. In general you should absolutely resist the requirement to proceed according to the RCM JA1011 standard, even if you do wish for a process that closely resembles RCM. The methodology recommended here was developed over a 15 year period by the nuclear industry, and has continued to be supported by the Nuclear Regulatory Commission and other regulators. The industry has sponsored studies that demonstrate results that are absolutely equivalent to those from using RCM. The methodology described here remains true to the RCM philosophy and can plausibly be said to conform to the JA1011 standard, even though some of the fine points are worth debating. As described above, the approach also has some technical merits that go (usefully) beyond what RCM requires. You would also be well served by assigning a senior engineer to study the alternative approaches and become thoroughly familiar with their pros and cons before starting your project.

Figure 1: PM Template Application