MAINTENANCE MANAGEMENT

Presented by

ALADON LTD

44 Regent Street, Lutterworth, Leicestershire
Phone: 01455 557542   Fax: 01455 553993

MAINTENANCE MANAGEMENT

21

CONTENTS | PAGE

1 Introduction 1

One of the most striking features of late twentieth century economic activity is the extent to which businesses have come to depend on physical assets. This has taken place in two stages. First, in industries ranging from earthmoving to high-precision assembly, machines have taken over jobs that people used to do. This led to a huge reduction in the numbers of people employed in industry. Most of those who remained became machine operators (and maintainers) rather than manual laborers. Next, automated control systems started taking over the role of the machine operators, to the extent that wholly automated processes are becoming increasingly commonplace. This applies to industries as diverse as manufacturing, telecommunications, utilities, warehousing and mass transport.

Apart from leading to massive improvements in productivity and consistency of output, the growth in mechanization and automation has three key implications.

Firstly, modern physical assets cost a great deal to acquire, so much so that the capital invested in plant and equipment is now the dominant feature of many balance sheets. In these cases, the primary measure of business performance - return on investment - is to a considerable extent measuring return on what is invested in physical assets. (In other words, physical assets represent a major part of the bottom line of the ROI equation.) In organizations which lease all or most of their equipment, the cost of servicing the leases also heavily influences the magnitude of the net wealth generated by the enterprise.

Secondly, modern physical assets cost a great deal to maintain, to the extent that in some industries, maintenance is now the second or third highest category of operating costs. In some cases, it is even the highest. As a result, in only forty years, maintenance has moved from almost nowhere to the top of the league as a cost control priority.

Finally and perhaps most important, the ability of highly mechanized and automated enterprises to satisfy their customers depends at least as much on the ongoing performance of their physical assets as it does on the performance of their people. This in turn demands physical assets that do what their users want at the moment they are put into service, and also that these assets continue to satisfy their users throughout their technologically useful lives.

In essence, these three issues represent three major managerial challenges:

- to minimise the cost of acquiring physical assets
- to minimise the cost of maintaining the assets
- to ensure that the assets continue to perform satisfactorily.

How well these three challenges are met profoundly influences the overall effectiveness of any enterprise that makes use of physical assets. As a result, the formulation and execution of suitable physical asset management strategies has become a top priority – if not a prerequisite for survival – in many industries. However, meeting these three challenges is not easy, for two reasons:

- contradictory actions often appear to be required to meet each of the challenges. For example, lowering acquisition costs (in other words, buying cheaper equipment) often increases maintenance costs. Similarly, if maintenance costs are cut inappropriately, equipment performance begins to suffer (often at a far greater cost to the enterprise). More so than in most other management disciplines, this means that successful physical asset management entails continuously striving to find a balance between apparently conflicting business priorities

- our understanding of what must be done to manage physical assets successfully – from both the technical and the organizational viewpoint – has changed profoundly in the recent past. The point has been reached where a great deal of accepted wisdom in this field is now known to be incorrect, often actively counterproductive and sometimes downright dangerous. Any organization which depends on physical assets and which hopes to develop a physical asset management strategy that will take it safely and cost-effectively into the twenty-first century must identify and eliminate these misconceptions.

The next part of this paper considers the most important areas of change. Thereafter, it reviews the role of the physical asset management function in the light of the changes, describes an asset management strategy formulation process that enables users to give effect to all the most important changes simultaneously, and ends with a look at the meaning of the term ‘responsible custodianship’.

2 A New Paradigm 2

The subject of change dominates nearly everything currently being written about management. All disciplines are being exhorted to adapt to changes in organization design, in technology, in leadership skills, in communication – in fact, in virtually every aspect of working life.

Perhaps nowhere is this felt more broadly and deeply than in the field of physical asset management.

A striking feature of this phenomenon is the number of changes which have occurred together. Some have occurred at a strategic – almost philosophical – level, while others are more tactical – or technical – in nature. Even more striking is the extent of the changes. Not only do they involve radical changes of direction, but a few ask us to come to terms with entirely new concepts.

This part of this paper identifies twelve key areas of change. Each of them *on its own* is sufficiently far-reaching to merit a great deal of attention in most organizations. Together they amount to a whole new paradigm. Accommodating this paradigm shift means that for many of its exponents, the management of physical assets is going to become a monumental exercise in change management over the next few years.

Each of the changes on its own is also sufficient to form the subject of one – if not several – books, so a short paper like this cannot hope to explore all the changes in detail. In fact, it goes to the opposite extreme by reducing each area of change to two maxims followed by a short explanation. In each case, one maxim attempts to summarize the way things used to be, while the other summarizes the way things are – or should be – now.

MAXIM 1

OLD

The primary objective of maintenance is to optimise plant availability at minimum cost

NEW

Maintenance affects all aspects of the business: safety, environmental integrity energy efficiency and product quality, not just plant availability and cost

Downtime has always affected the productive capability of physical assets by reducing output, increasing operating costs and interfering with customer service. By the 1960's and 1970's, this was already a major concern in the mining, manufacturing and transport sectors. In manufacturing, the effects of downtime are being aggravated by the worldwide move towards just-in-time systems, where reduced stocks of work-in-progress mean that quite small breakdowns are now much more likely to stop a whole plant. As a result, the lost output associated with such failures can cost hundreds or even thousands of times as much as they cost to repair. In recent times, the growth of mechanization and automation means that reliability and availability have also become key issues in sectors as diverse as health care, data processing, telecommunications, banking and building management.

The cost of maintenance has also been rising at a steady pace over the past few decades, in absolute terms and as a proportion of total expenditure, so reducing – or at least containing – these costs is indeed a high priority.

The importance of these two aspects of asset management means that many maintenance managers still tend to view them as the only significant objectives of maintenance.

However, this is no longer the case, because the maintenance function now has a wide range of additional objectives. These are summarized in the following paragraphs.

Greater automation means that more and more failures affect our ability to achieve and sustain satisfactory quality standards. This applies as much to *standards of service* as it does to *product quality*. For instance equipment failures affect gas, water and electricity distribution and the punctuality of transport networks as much as they interfere with the consistent achievement of specified tolerances in manufacturing.

Safety is also becoming a dominant issue in the world of maintenance. Equipment failure has played a part in most of the worst accidents in industrial history - accidents which have become bywords, such as Bhopal, Chernobyl and Piper Alpha. In fact, the extent to which the physical health of most organizations now depends on the continued physical and functional integrity of their assets means that the pressure upon maintainers to manage their assets in the most responsible fashion possible is becoming extraordinarily intense. Not only is this pressure arising from the expectations of our employees and our customers, but it is attracting the attention of regulators. Government bodies like OSHA, the FDA, the FAA and the EPA in the USA and the HSE in the UK, in addition to regional and municipal regulatory bodies, are not only demanding much greater precision and clarity in our asset management policies, but they are also asking us to be able to prove that what we are doing is sensible and defensible. The sanctions which they apply if we are thought to have got it wrong are also becoming steadily more ferocious. For example, the British government has recently started prosecuting senior executives of organizations for a new class of crime called 'corporate manslaughter', in cases where fatalities are thought to be the result of irresponsible custodianship of physical assets.

Another result of growing automation is the rising number of failures which have serious environmental consequences, at a time when standards in this area are rising fast. Many parts of the world are reaching the point where organizations either conform to society's *environmental expectations*, or they get shut down. This adds an order of magnitude to our dependence on the integrity of our physical assets – one that goes beyond cost and becomes a simple matter of organizational survival.

At the same time as our dependence on physical assets is growing, so too is their *cost – to operate* and *to own.* To secure the maximum return on the investment which they represent, they must be kept working efficiently for as long as their users want them to.

These developments mean that physical asset management now plays an increasingly central role in preserving *all* aspects of the physical, financial and competitive health of organizations that rely on such assets. If they wish to survive and prosper, these organizations need to equip themselves with the tools required to address these issues continuously, proactively and directly, rather than deal with them on an ad hoc basis when time permits

MAXIM 2

OLD

Maintenance is all about preserving physical assets

NEW

Maintenance is all about preserving the functions of physical assets

Most people become engineers because they feel at least some affinity for things, be they mechanical, electrical or structural. This leads us to derive pleasure from assets in good condition, but feel offended by assets in poor condition.

These reflexes have always been at the heart of the concept of preventive maintenance. They have given rise to concepts like “asset care”, which as the name implies, seeks to care for assets *per se*. They have also led maintenance strategists to believe that maintenance is all about preserving the inherent reliability or built-in capability of any asset.

In fact, this is not so. As we gain a deeper understanding of the role of assets in business, we begin to appreciate the significance of the fact that any physical asset is put into service because someone wants it to do something. So it follows that when we maintain an asset, *the state which we wish to preserve must be one in which it continues to do whatever its users want it to do*. This in turn implies that we have to focus our attention on maintaining what each asset *does* rather than on what it *is*.

Clearly, before we can do this, we must gain a crystal clear understanding of the functions of each asset together with the associated performance standards

For instance, Figure 1 shows a pump with a nominal capacity of 400 liters/minute pumping water into a tank from which it is being drawn at a rate of 300 liters/minute. In this case, the primary function of the pump is “to supply water to the tank at not less than 300 liters/minute”. Any maintenance program for the pump should try to ensure that its performance does not drop below 300 liters/minute. (In seeking to ensure that the tank does not run dry, the maintenance program does *not* try to ensure that the pump continues “to supply water to the tank at 400 liters/minute”.)

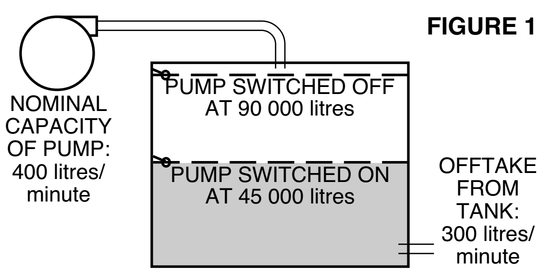

FIGURE 1

However, if the offtake from the tank is changed to 350 liters/minute, the primary function is changed accordingly, so the maintenance program must also change to accommodate the higher performance expectation.

Note that there should always be a gap between the nominal capacity of the pump in the example (400 liters/minute) and offtake from the tank (300 liters/minute). The gap is needed because the laws of physics tell us that any organized system that is exposed to the real world will deteriorate. The end result of this deterioration is total disorganization (also known as ‘chaos’ or ‘entropy’), unless steps are taken to arrest whatever process is causing the system to deteriorate.

So if deterioration is inevitable, it must be allowed for. This means that when any asset is put into service, it must be able to deliver *more* than the minimum standard of performance desired by the user. What the asset is able to deliver is known as its *initial capability* (or inherent reliability). Figure 2 illustrates the right relationship between this capability and desired performance.

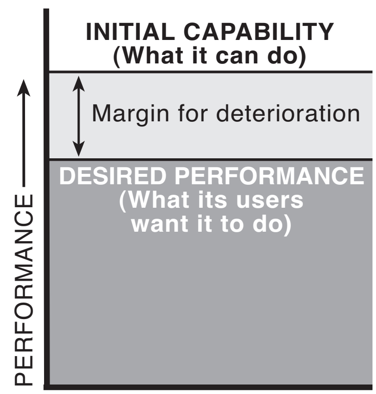

FIGURE 2

Functions and performance expectations not only cover output. They also concern issues such as product quality, customer service, economy and efficiency of operation, control, containment, comfort, protection, compliance with environmental regulations, structural integrity and even the physical appearance of the asset.

MAXIM 3

OLD

Most equipment becomes more likely to fail as it gets older

NEW

Most failures are not more likely to occur as equipment gets older

For decades, conventional wisdom suggested that the best way to optimize the performance of physical assets was to overhaul or replace them at fixed intervals. This was based on the premise that there is a direct relationship between the amount of time (or number of cycles) equipment spends in service and the likelihood that it will fail, as shown in Figure 3. This suggests that most items can be expected to operate reliably for a period "X", and then wear out.

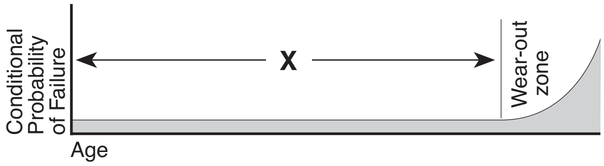

FIGURE 3

Classical thinking held that X could be determined from historical records about equipment failure, enabling users to take preventive action shortly before the item is due to fail in future. This predictable relationship between age and failure is indeed true for some failure modes. It tends to be found where equipment comes into direct contact with the product. Examples include pump impellers, furnace refractories, machine tooling, screw conveyors, valve seats, crusher liners and so on. Age-related failures are also often associated with fatigue and corrosion.

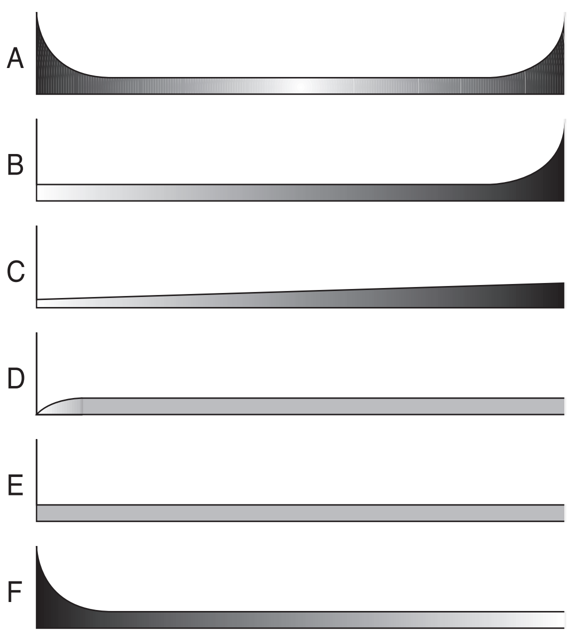

FIGURE 4

However, equipment in general is much more complex than it was even fifteen years ago. This has led to startling changes in the patterns of equipment failure, as shown in Figure 4. The graphs show conditional probability of failure against operating age for a wide variety of failure modes that can affect electrical, mechanical and structural items.

Pattern A is the well-known bathtub curve, and pattern B is the same as Figure 3. Pattern C shows slowly increasing probability of failure with no specific wear-out age. Pattern D shows low failure probability to begin with then a rapid increase to a constant level, while pattern E shows a constant probability of failure at all ages Pattern F starts with high infant mortality and drops eventually to a constant or very slowly decreasing failure probability.

Studies performed on several hundred mechanical, electrical and structural components of civil aircraft showed that 4% of the items studied conformed to pattern A, 2% to B, 5% to C, 7% to D, 14% to E and no fewer than 68% to pattern F. (The distribution of these patterns in aircraft is not necessarily the same as in industry, but as equipment grows more complex, more and more items conform to patterns E and F.)

These findings contradict the belief that there is always a connection between reliability and operating age – the belief which led to the idea that the more often an item is overhauled, the less likely it is to fail. In practice, this is hardly ever true. Unless there is a dominant age-related failure mode, fixed interval overhauls or replacements do little or nothing to improve the reliability of complex items.

Most maintenance professionals are aware of these findings, and coming to terms with the reality of randomness after decades in the bathtub. However, the fact that the bathtub curve still features in so many texts on maintenance is testimony to the almost mystical faith that some people still place in the relationship between age and failure. In practice, this faith has two serious drawbacks, as follows:

- it leads to the belief that if we don't have any hard evidence at all about the existence of an age-related failure mode, it is wise to overhaul the item anyway from time to time "just-in-case" such a failure mode does exist. This ignores the fact that overhauls are extraordinarily invasive undertakings that massively upset stable systems. As such, they are highly likely to induce infant mortality, and so cause the very failures that they seek to prevent. This is illustrated in Figure 5.

FIGURE 5

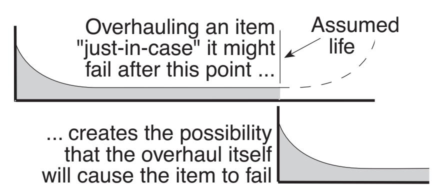

- at a more philosophical level, bathtub believers convince themselves that it is more conservative (in other words, safer) to assume that everything has a life – and hence to overhaul equipment on the basis of an assumed life – than to assume it could fail at random. After implementing overhaul programs based on this assumption, they then assume that no failures should occur between overhauls, and that any that do occur cannot be attributed to maintenance because the item has just been overhauled. The possibility that the overhaul itself may be the cause of the failure is usually completely lost on such people. More seriously, they simply refuse to accept the most important conclusion associated with maxim 4, which is *that in the absence of any evidence to the contrary, it is more conservative to assume that any failure can occur at* ***any*** *time (in other words, at random), and not to assume that it will only occur after some fixed amount of time in service.*

MAXIM 4

OLD

Proactive maintenance is all about preventing failures

NEW

Proactive maintenance is about avoiding, eliminating or minimizing the consequences of failures

A detailed analysis of an average industrial undertaking is likely to yield between five and ten thousand possible ways equipment can fail (or *failure modes*). Each of these failures will affect the organization in some way, but in each case, the effects will be different. They might affect operations. They might also affect product quality, customer service, safety or the environment. They will all take time and cost money to repair.

It is these consequences which most strongly influence the extent to which we try to prevent each failure. If a failure mode has serious consequences, we are likely to go to great lengths to try to prevent it. If it has little or no effect, then we may decide to undertake no preventive action.

In other words, the consequences of failures are far more important than their technical characteristics.

For example, one failure that could affect the pump shown in Figure 1 is “bearing seizes due to normal wear and tear”. If it takes 4.5 hours to replace a failed bearing, and if the unanticipated failure of the bearing only comes to the attention of the operators when the level in the tank drops to the low level switch. At this point, the tank only contains a 2.5 hour supply of water, so it would run dry and remain empty for 2 hours while the bearing is under repair.

One condition-based task which could apply to this bearing is to monitor vibration levels. If incipient failure is detected, the first priority of the operators would be to fill the tank before the bearing seizes, thus giving themselves 5 hours to do a 4 hour job. This in turn enables them to avoid the consequences of an empty tank. *The task does not "save" the bearing* – that is doomed whatever happens.

Another example of the extent to which failure consequences influence the selection of maintenance policies is shown by three otherwise identical pumps in Figure 6. Pump A stands alone, so if it fails, operations will be affected sooner or later. As a result the users and/or maintainers of Pump A are likely to make some effort *to anticipate or prevent its failure*. (How hard they try will be governed both by the effect on operations and by the severity and frequency of the failures of the pump.)

However, if pump B fails, the operators simply switch to pump C, so the only consequence of the failure of pump B is that it must be repaired. As a result, it is likely that the operators of B would at least consider letting it *run to failure* (especially if the failure of B does not cause significant secondary damage.)

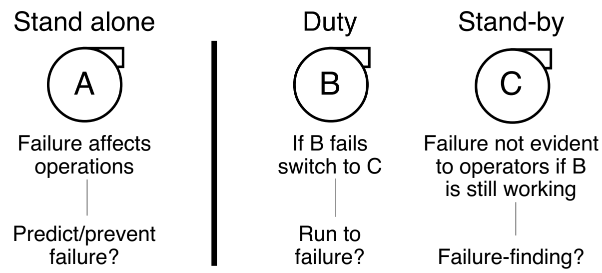

FIGURE 6

On the other hand, if pump C fails while pump B is still working (for instance if someone cannibalizes a part from C), it is likely that the operators will not even know that C has failed unless or until B also fails. To guard against this possibility, a sensible maintenance strategy might be *to run C from time to time to find out whether it has failed.*

This example shows how three identical assets can have three totally different maintenance policies, because the failure consequences are different in each case.

Both examples show that the main reason for doing any kind of proactive maintenance is to avoid, reduce or eliminate the *consequences* of failure. A formal review of failure consequences focuses attention on maintenance tasks that have most effect on the performance of the organization, and diverts energy away from those that have little or no effect. This helps ensure that whatever is spent on maintenance is spent where it will do the most good.

MAXIM 5

OLD

Generic maintenance programs can be developed for most types of physical assets

NEW

Geenric maintenance programs only apply to equipment with the same operating context, functions and performance standards

The belief that generic maintenance policies can and should be applied to most types of assets lies at the heart of nearly all traditional maintenance programs. For instance, one often hears people say things like "the maintenance policy which we apply to all our pumps is X" or "we have a type Y calibration policy for all our instruments".

However, scientific maintenance strategy formulation techniques reveal that the inappropriate use of generic maintenance is one of the main reasons why so many traditional maintenance programs fail to achieve their full potential. The following paragraphs explain why generic policies should be treated with great caution:

- *functions:* the narrative accompanying Figure 1 on page 2 explained how different performance standards applied to otherwise identical assets call for different standards of maintenance. (This is especially true where otherwise identical machines are used to produce products which have widely differing quality standards.)
- *failure modes:* when otherwise identical equipment is used in even slightly different locations (an area of high humidity, an unusually dusty environment) or to perform slightly different tasks (cutting a harder than usual metal, operating at a higher temperature, pumping a more abrasive or a more acidic liquid), the possible failure modes vary drastically. This in turn means that failure management strategies need to vary accordingly
- *failure consequences:* The example of the three pumps in Figure 6 showed how three assets that would traditionally be seen as natural candidates for generic maintenance could in fact be subjected to three totally different maintenance programs, because the failure consequences are different in each case.
- *maintenance tasks:* different organizations – or even different parts of the same organization – seldom employ people with identical skillsets. This means that people working on one asset may prefer to use one type of proactive technology (say high-tech condition monitoring), while another group working on an identical asset may be more comfortable using another (say a combination of performance monitoring and the human senses). It is surprising how often this does not matter, as long as the techniques chosen are cost-effective. In fact, many maintenance organizations are starting to realize that there is often more to be gained from ensuring that the people doing the work are comfortable with what they are doing than it is to compel everyone to do the same thing. (The validity of different tasks is also affected by the operating context of each asset. For instance, think how background noise levels affect checks for noise.)

All of this means that special care must be taken to ensure that the operating context, functions and desired standards of performance are all virtually identical before applying a maintenance policy designed for one asset to another.

MAXIM 6

OLD

Comprehensive data about failure rates must be available before it is possible to develop successful maintenance strategies

NEW

Decisions about maintenance will nearly always have to be made with inadequate hard data about failure rates

A surprising number of people believe that effective maintenance policies can only be formulated on the basis of extensive historical information about failure. Thousands of manual and computerized technical history recording systems have been installed around the world on the basis of this belief. It has also led to great emphasis being placed on the failure patterns discussed earlier in this paper. Yet from the maintenance viewpoint, these patterns are fraught with practical difficulties, conundrums and contradictions. Some of these are summarized below:

- *Sample size and evolution:* Large industrial processes often have only one or two assets of any one type. They tend to be brought into operation in series rather than simultaneously. This means that sample sizes tend to be too small for statistical procedures to carry much conviction.
- *Complexity:* The number and diversity of assets present in most industrial undertakings means that it is simply not possible to develop a complete analytical description of the reliability characteristics of an entire undertaking – or even any major asset within the undertaking.
- *Reporting failure:* Further complications arise due to differences in reporting policy from one organization to another. For example, an item may be replaced on one site because it is failing while on another it is replaced because it has failed. Similar differences are caused by different performance expectations. So what is a failure in one organization - or even one part of an organization - might not be a failure in another. This can result in two different sets of failure data for two apparently identical items.
- *The ultimate contradiction:* The whole question of technical history is bedevilled by the fact that if we are collecting data about failures, it must be because we are not preventing them. The implications of this are summed up most clearly by Resnikoff (1978) in the following statement:

  "The acquisition of the information thought to be most needed by maintenance policy designers – information about critical failures – is in principle unacceptable and is evidence of the failure of the maintenance program. This is because critical failures entail potential (in some cases, certain) loss of life, but there is no rate of loss of life which is acceptable to (any) organization as the price of failure information to be used for designing a maintenance policy. Thus the maintenance policy designer is faced with the problem of creating a maintenance system for which the expected loss of life will be less than one over the planned operational lifetime of the asset. This means that, both in practice and in principle, the policy must be designed without using experiential data which will arise from the failures which the policy is meant to avoid."

  If despite our best efforts, a critical failure actually does occur, Nowlan and Heap (1978) go on to make the following comments about the role of actuarial analysis:

  "The development of an age-reliability relationship, as expressed by a curve representing the conditional probability of failure, requires a considerable amount of data.

When the failure is one which has serious consequences, this body of data will not exist, since preventive measures must of necessity be taken after the first failure. Thus actuarial analysis cannot be used to establish the age limits of greatest concern - those necessary to protect operating safety."

This brings us to the ultimate contradiction concerning the prevention of failures with serious consequences and historical information about such failures: that successful proactive maintenance entails preventing the collection of the historical data which we think we need in order to decide what proactive maintenance we ought to be doing.

This contradiction applies in reverse at the other end of the scale of consequences. Failures with minor consequences tend to be allowed to occur precisely because they do not matter very much. As a result, large quantities of historical data are available concerning these failures, which means that there will be sufficient material for accurate actuarial analyses. These may even reveal some age limits. However, because the failures don't matter much, it is highly unlikely that the resulting fixed interval maintenance tasks will be cost effective. So while the actuarial analysis of this information may be precise, it is also likely to be a waste of time.

Perhaps the most important conclusion to arise from the above comments is that maintenance professionals should turn their attention away from counting failures (in the hope that an elegantly constructed scorecard will tell us how to play the game in the future), towards anticipating or preventing failures that matter.

So to be truly effective, we simply have to get comfortable with the idea of uncertainty, and deploy strategies which enable us to deal with it confidently. We also need to recognize that if the consequences of too much uncertainty cannot be tolerated, then we must change the consequences. In extreme cases of uncertainty, the only way to do so may be to abandon the process concerned.

MAXIM 7

OLD

The probability of catastrophic failures can be almost eliminated by fitting suitable protection

NEW

Protection can also fail, so the risks associated with protected systems still need to be managed

The growth in the number of ways in which equipment can fail has led to corresponding growth in the variety and severity of failure consequences. Protective devices are being used increasingly in an attempt to eliminate (or at least reduce) these consequences. These devices work in one of five ways:

- to alert operators to abnormal conditions
- to shut down the equipment in the event of a failure
- to eliminate or relieve abnormal conditions which follow a failure and which might otherwise cause much more serious damage
- to take over from a function which has failed
- to prevent dangerous situations from arising.

In essence, the function of these devices is to ensure that the consequences of the failure of the protected function are much less serious than they would be if there were no protection. So any protective device is in fact part of a system with at least two components:

- the protective device
- the protected function.

One unfortunate feature of such systems is that many of their users think that once a protective device has been fitted, no further action needs to be taken to manage the associated risks. As a result (and as discussed further in the next section of this paper), many of these devices receive no attention at all. However, it has now become painfully apparent that protective devices can also fail. This creates the possibility that the protected function could fail while the protective device is in a failed state. This in turn is known as a multiple failure, as shown in Figure 7.

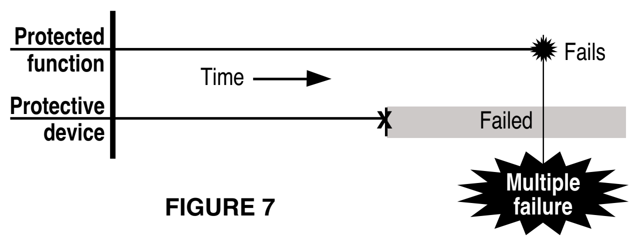

FIGURE 7

In practice, it is possible to vary both the probability of failure of a protected function and (especially) the downtime of the protective device by adopting suitable maintenance and operating policies. As a result, it is also possible to reduce the probability of multiple failures to almost any desired level within reason by adopting such policies. (Zero is of course an unattainable ideal.)

The probability that is considered tolerable for any multiple failure depends on its consequences. Sometimes levels of tolerability are specified by regulatory authorities, but in the vast majority of cases the assessment has to be made by the users of the asset. Since these consequences vary hugely from system to system, what is deemed to be tolerable varies equally widely. This means that there no universal standards of risk that can be applied to all systems of a particular type (at least, not yet).

But *someone* has to make a decision as to what level of risk is tolerable before it is possible to decide what must be done to design, operate and maintain protected systems. In fact, coming to terms with the fact that this is a manageable variable *which must therefore be managed* is currently one of the biggest challenges facing the users of physical assets.

MAXIM 8

OLD

There are three basic types of maintenance: predictive, preventive and corrective

NEW

There are four basic types of maintenance: predictive, preventive, detective and corrective

Most of what has been written to date on the general subject of maintenance strategy refers to three – and only three – types of maintenance: predictive, preventive and corrective.

Predictive (or condition-based) tasks entail checking if something is failing. Preventive maintenance usually means overhauling items or replacing components at fixed intervals. Corrective maintenance means fixing things either when they are found to be failing or when they have failed.

However, there is a whole family of maintenance tasks which falls into none of the above categories.

For example, when we periodically activate a fire alarm, we are not checking if it is failing. We are also certainly not overhauling or replacing it, and nor are we repairing it.

We are simply checking if it still works.

Tasks designed to check whether something still works are known as "functional checks" or "failure-finding tasks". (In order to rhyme with the other three families of tasks, they are sometimes called "detective" tasks because they are used to detect if something has failed.)

Detective maintenance or failure-finding applies only to hidden or unrevealed failures, and hidden failures in turn only affect protective devices.

If one applies scientific maintenance strategy formulation techniques to almost any modern, complex industrial system, one finds that up to 40% of failure modes fall into the hidden category. Furthermore, up to 80% of these failure modes require failure-finding, so typically, *one third of the tasks generated by comprehensive, correctly applied maintenance strategy development programs are detective tasks.*

On the other hand, the same analytical techniques reveal that it is not unusual for condition monitoring to be technically feasible for no more than 20% of failure modes, and worth the investment in less than half these cases. (This is not meant to imply that condition monitoring should be not be used – where it is good it is very, very good – but that we must also remember to develop appropriate strategies for managing the other 90% of our failure modes.)

A rather more troubling finding is that most traditionally derived maintenance programs provide for fewer than one third of protective devices to receive any attention at all (and then usually at inappropriate intervals). The people who operate and maintain the plant covered by these traditional programs are aware that another third of these devices exist but pay them no attention, while it is not unusual to find that no-one even knows that the final third exist. This lack of awareness and attention means that most of the protective devices in industry – our last line of protection when things go wrong – are maintained poorly or not at all.

This situation is completely untenable.

If industry is serious about safety and environmental integrity, then the whole question of detective maintenance – failure-finding – needs to be given top priority as a matter of urgency. As more and more maintenance professionals become aware of the importance of this neglected area of maintenance, it is likely to become a bigger maintenance strategy issue in the next decade than predictive maintenance has been in the last ten years.

MAXIM 9

OLD

The frequency of predictive tasks should be based on the frequency of the failure and/or the criticality of the item

NEW

The frequency of predictive tasks should be based on the failure development period (or "lead time to failure" or "P-F interval")

When people are discussing the frequency of predictive (or condition-based) maintenance tasks, one often hears either – sometimes both – of the following statements:

- it doesn't fail so often, so we don't need to check it so often
- we need to check more critical plant more often than less critical plant.

In both cases, the speakers are wrong. The frequency of predictive maintenance tasks has nothing to do with the frequency of failure and or with the criticality of the item. The frequency of any form of condition-based maintenance is based on the fact that most failures do not occur instantaneously, and that it is often possible to detect the fact that the failure is occurring during the final stages of deterioration.

Figure 8 shows this general process. It is called the P-F curve, because it shows how a failure starts and deteriorates to

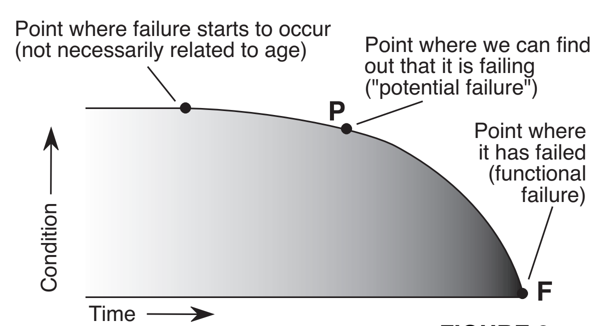

FIGURE 8

the point at which it can be detected (the potential failure point "P"). Thereafter, if it is not detected and suitable action taken, it continues to deteriorate – usually at an accelerating rate – until it reaches the point of functional failure ("F").

The amount of time (or the number of stress cycles) which elapse between the point where a potential failure occurs and the point where it deteriorates into a functional failure is known as the P-F interval, as shown in Figure 9.

FIGURE 9

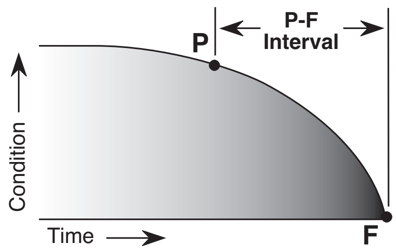

The P-F interval governs the frequency with which the predictive task must be done. The checking interval must be significantly less than the P-F interval if we wish to detect the potential failure before it becomes a functional failure.

The P-F interval can be measured in any units relating to exposure to stress (running time, units of output, stop-start cycles, etc), but it is most often measured in terms of elapsed time. For different failure modes, the P-F interval can vary from fractions of a second to several decades.

The amount of time needed to respond to any potential failures that are discovered also influences condition-based task intervals. In general, these responses consist of any or all of the following:

- take action to avoid the consequences of the failure
- plan corrective action so that it can be done without disrupting production and/or other maintenance activities
- organize the resources needed to rectify the failure.

The amount of time needed for these responses also varies, from a matter of hours (say until the end of an operating cycle or the end of a shift), minutes (to clear people from a building which is falling down) or even seconds (to shut down a machine or process which is running out of control) to weeks or even months (say until a major shutdown).

Unless there is a good reason to do otherwise, it is usually sufficient to select a checking interval equal to half the P-F interval. This ensures that the task will detect the potential failure before the functional failure occurs, while providing a net interval of at least half the P-F interval to do something about it. However, it is sometimes necessary to select a checking interval which is some other fraction of the P-F interval. For instance, Figure 10 shows how a P-F interval of 9 months and a checking interval of 1 month give a net P-F interval of 8 months.

FIGURE 10

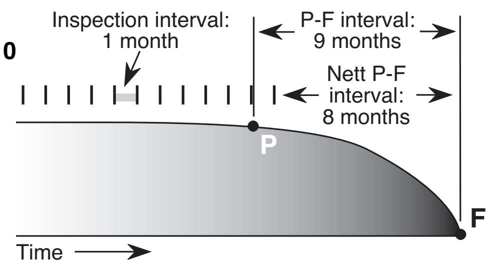

If the P-F interval is too short for it to be practical to check for the potential failure, or if the net P-F interval is too short for any sensible action to be taken once a potential failure is discovered, then the condition-based task is not appropriate for the failure mode under consideration.

MAXIM 10

OLD

Maintenance policies should be formulated by managers and maintenance schedules drawn up by suitably qualified specialists or external contractors (a top-down approach)

NEW

Maintenance policies should be formulated by the people closest to the assets. Management should provide the tools to help them make the right decisions, and ensure that the decisions are sensible and defensible

The compilation of maintenance schedules used to be the responsibility of the planned maintenance department. Planners often devoted immense amounts of time and energy to this exercise. However, more often than not, their schedules died when they reached the shop floor. This happened for two main reasons:

- *technical validity:* the planners who wrote the schedules were usually out of touch with the equipment (if they had ever been in touch to start with). As a result, they often had a less than adequate understanding of the functions, the failure modes and effects and the failure consequences of the assets for which their schedules were being written. This meant that the schedules were usually generic in nature, so people who were supposed to do them often saw them as being incorrect if not totally irrelevant
- *ownership:* people on the shop floor (supervisors and craftsmen) tended to view the schedules as unwelcome paperwork which appeared from some ivory tower and disappeared after it was signed off. Many of them learned that it was more comfortable just to sign off the schedules and send them back than it was to attempt to do them. (This led to inflated schedule completion rates which at least kept the planners happy.) The main reason for the lack of interest was undoubtedly sheer lack of ownership.

The only way around the problems of technical invalidity and lack of ownership is to involve shop floor people directly in the maintenance strategy formulation process. This is because they are the ones who really understand how the equipment works, what goes wrong with it, how much each failure matters and what must be done to fix it.

The best way to access their knowledge on a systematic basis is to arrange for them to participate formally in a series of meetings. However, in order to ensure that these meetings do not just become another bunch of inconclusive talkfests, the participants must first be trained in the most focused and effective strategy formulation techniques, and provided with skilled guidance in the application of these techniques.

Done correctly, this not only produces schedules with a much higher degree of technical validity than anything that has gone before, but it also produces an exceptionally high level of ownership of the final results.

(A word of caution at this stage: It is wise to steer clear of the temptation to use external contractors to formulate maintenance strategies. An outsider's sheer ignorance of almost all the issues discussed in connection with maxims 1 through 9 insofar as they affect *your* plant means that all you are likely to get is a set of elegantly completed forms that amount to little or nothing. Using such people to develop maintenance programs is to wander into the hazy – and dangerous – region where delegation becomes abdication.)

MAXIM 11

OLD

The maintenance department on its own can develop a successful, lasting maintenance programme

NEW

A successful, lasting maintenance program can only be developed by maintainers and users of the assets working together

Maxim 10 above reminds us of the need to involve shop floor people as well as managers in the maintenance strategy development process. Maxim 11 concerns what is often a much more difficult challenge in many organizations – the almost impenetrable divide between the maintenance and operations functions.

In fact, as the second maxim in this series makes clear, maintenance is all about ensuring that assets continue to function to standards of performance required by the users. In most cases, the "users" are the production or operations functions. This means that modern maintenance strategy formulation starts by asking the users what they want, with a view to setting up asset management programs whose sole objective is to ensure that the users get what they want. Clearly, for this to be possible, the users must be prepared to specify exactly what they require. (If they do not bother to state the performance they require from each asset with adequate precision, then of course they cannot hold maintenance responsible for delivering that performance.) Both users and maintainers must also take care at this stage to satisfy themselves that the asset is capable of delivering the required performance to begin with.

In addition to spelling out what they want the asset to do, operators also have a vital contribution to make to the rest of the strategy formulation process.

By participating in a suitably focused FMEA, they learn a great deal about failure modes caused by human error, and hence what they must do to stop breaking their machines. They also play a key role in evaluating failure consequences, and they have invaluable personal experience of many of the most common warnings of failure (especially those detected by the human senses). Finally, involvement in this process helps users to understand much more clearly why they sometimes need to release machines for maintenance, and also why operators need to be asked to carry out certain maintenance tasks.

In short, from a purely technical point of view, it is rapidly becoming apparent that it is virtually impossible to set up a viable, lasting maintenance program in most industrial undertakings without involving the users of the assets. If their involvement can be secured at all stages in the process, that notorious barrier rapidly starts to disappear and the two departments start to function, often for the first time ever, as a genuine team.

MAXIM 12

OLD

Equipment manufacturers are in the best position to develop maintenance programs for new physical assets

NEW

Equipment manufacturers can only play a limited (but still important) role in developing maintenance programs for new assets

A universal feature of traditional asset procurement is the insistence that the equipment manufacturer should provide a maintenance program as part of the supply contract for new equipment. Apart from any thing else, this implies that manufacturers know everything that needs to be known to draw up suitable maintenance programs.

In fact, manufacturers are usually at best no better informed than traditional maintenance planners about the operating context of the equipment, desired standards of performance, context-specific failure modes and effects, failure consequences and the skills of the user's operators and maintainers. More often the manufacturers know nothing at all about these issues. As a result, schedules compiled by manufacturers are nearly always generic, with all the drawbacks discussed under maxim 5.

Equipment manufacturers also have other agendas when specifying maintenance programs (not least of which is to sell spares). What is more, they are either committing the users' resources to doing the maintenance (in which case they don't have to pay for it, so they have little interest in minimizing it) or they may even be bidding to do the maintenance themselves (in which case they have a vested interest in doing as much as possible).

This combination of extraneous commercial agendas and ignorance about the operating context means that maintenance programs specified by manufacturers tend to embody a high level of over-maintenance (sometimes ludicrously so) coupled with massive over-provisioning of spares. Most maintenance professionals are aware of this problem. However, despite our awareness, most of us persist in demanding that manufacturers provide these programs, and then go on to accept that they must be followed in order for warranties to remain valid (and so bind ourselves contractually to doing the work, at least for the duration of the warranty period).

None of this is meant to suggest that manufacturers mislead us deliberately when they put together their recommendations. In fact, they usually do their best in the context of their own business objectives and with the information at their disposal. If anyone is at fault, it is really us – the users – for making unreasonable requests of organizations which are not in the best position to fulfil them.

A growing number of users solve this problem by adopting a completely different approach to the development of maintenance programs for new assets. This entails asking the manufacturer to supply experienced field technicians to work alongside the people who will eventually operate and maintain the equipment, to develop programs which are satisfactory to both parties.

When adopting this approach, issues such as warranties, copyrights, languages which the participants should be able to speak fluently, technical support, confidentiality, and so on should be handled at the request for proposal/contracting stage, so that everyone knows what to expect of each other.

Note the suggestion to use field technicians rather than designers (designers are often surprisingly reluctant to admit that their designs can fail, which reduces their ability to help develop a sensible failure-management program). The field technicians should of course have unrestricted access to specialist support to help them answer difficult questions.

In this way, the *user* gains access to the most useful information that the manufacturer can provide, while still developing a maintenance program which is most directly suited to the context in which the equipment will actually be used. The *manufacturer* may lose a little in up-front sales of spares and maintenance, but will gain all the long-term benefits associated with improved equipment performance, lower through-life costs and a much better understanding of the real needs of his customer. A classic win-win situation.

CONCLUSION

OLD **O** — **NEW** **X**

If one takes a moment to review the breadth and depth of the paradigm shift implicit in the foregoing paragraphs, it soon becomes apparent just how far most organizations have to move in order to adopt the new maxims. It simply cannot happen overnight.

In fact, organizations that seek an effective, enduring maintenance program that has universal support should not lose sight of the fact that improvement is a journey, not a destination (the essence of the Kaizan philosophy). The first step on this journey is to change the way people think.

At this point in time, if anyone involved in the operation or maintenance of physical assets is asked what they think about the old maxims listed above, the vast majority will answer more-or-less as shown by the O's in Figure 11. In fact, the correct answers are generally as shown by the X's.

If people still think as shown by the O's, they will not change their behavior and the organization will continue to live – dangerously – in the past. So as mentioned above, if any organization wishes to apply the new maxims, the first step is to change the way its people think. The second step is to get them to apply their changed thought processes – systematically, sensibly and defensibly – to the physical assets under their control.

The next sections of this paper starts considering how this can be done by proposing a maintenance mission statement that reflects the new paradigm and goes on to describe an associated strategy formulation process.

	True	Partly True	Largely False	False
1	O			X
2	O			X
3		O	X	
4		O		X
5		O		X
6	O			X
7	O			X
8			O	X
9		O		X
10		O	X	
11		O		X
12	O			X

FIGURE 11

3 A Maintenance Mission 3

The number and the extent of the changes outlined in Part 2 of this paper – all of which must be dealt with urgently – mean that it is easy for anyone who wishes to implement the changes to get lost. In this environment, just as most major corporations develop formal mission statements to help them maintain a steady course in a changing world, it is worth developing a formal mission statement to help maintenance do likewise.

Perhaps a good place to start would be Maxim 2 above, which recognizes that every physical asset is put into service to fulfil a specific function or functions. This suggested that when we maintain an asset, the state we wish to preserve must be one in which it continues to do whatever its users want it to do. This shift in emphasis – from preserving what each asset *is* to preserving what it *does* – should be acknowledged in the mission statement.

The mission statement must also recognize the 'customers' of the maintenance service. Maintainers serve three distinct sets of customers – the owners of the assets, the users of the assets (usually the operators), and society as a whole. Owners are satisfied if their assets generate a satisfactory return on the investment made to acquire them. Users are satisfied if each asset continues to do whatever they want it to do to a standard of performance which they consider to be satisfactory. Society as a whole is satisfied if assets do not fail in ways which threaten public safety or the environment.

If things didn't fail they wouldn't need maintenance. So the technology of maintenance is all about finding and applying suitable ways of managing failure. Failure management techniques include predictive, preventive, detective and corrective maintenance, run to failure and one-time changes to the design of the asset or the way it is operated.

Each category includes a host of options, some more effective than others. Maintainers not only need to learn what these options are, but they also have to decide which are worthwhile in their own organizations. If they make the right choices, it is possible to improve asset performance and *at the same time* contain and even reduce the cost of maintenance. If they make the wrong choices, new problems are created while existing problems get worse. So the mission statement should emphasize the need to make the most cost-effective choices from the full array of options.

When considering failure management options, Maxim 4 reminds us that failures only attract attention because they have consequences. The severity and frequency of the consequences incurred by a failure dictates whether any failure management technique is worth applying. So the mission statement should acknowledge the key role of consequence avoidance in maintenance.

It should also acknowledge that most of us work in a highly resource constrained environment. The most efficient maintainers are those who apply the resources that they do need – people, spares and tools – as cost-effectively as possible, but not so cheaply that they damage the long-term functionality of their assets. In other words, the cost of ownership of the assets must be minimized throughout their useful lives, not just to the end of the next accounting period.

Finally, the mission statement must recognize that maintenance depends on people – not only maintainers, but also operators, designers and vendors. So it should acknowledge the need for everyone involved with the assets to share a common and correct understanding of what needs to be done, and to be able and willing to do whatever is needed right first time every time. All this suggests the following as a possible maintenance mission statement:

To preserve the functions of our physical assets throughout their technologically useful lives
- ***to the satisfaction of their owners, of their users and of society as a whole***
- ***by selecting and applying the most cost-effective techniques for managing failures and their consequences***
- ***with the active support of all the people involved.***

4 Formulating a Maintenance Strategy 4

It is one thing to decide on a mission. It is quite another to develop and implement a strategy that enables the maintenance enterprise to accomplish that mission.

Given all the day-to-day pressures faced by maintenance managers, the first question is where does one start? Buy a new maintenance management system (MMS)? Reorganize? Invest in truckloads of condition monitoring equipment? Knock the whole place down and build a new one?

The answer lies at the beginning of the mission statement, which states that our mission is *to preserve the functions of our assets*. It is only when these functions have been defined that it becomes clear exactly what maintenance is trying to achieve, and also precisely what is meant by "failed". This makes it possible to move on to the next step, which is to identify the reasonably likely causes and effects of each failed state.

Once failure causes (or failure modes) and effects have been identified, we are then in a position to assess how and how much each failure matters. This in turn enables us to determine which of the full array of failure management options should be used to manage each failure mode.

When the tasks that need to be done - the maintenance requirements of each asset - have been clearly identified, the next step is to decide sensibly what resources are needed to do each task. "Resources" consist of people and things, so the following questions must now be answered:

- who is to do each task: a skilled maintainer? the operator? a contractor? the training department (if training is required)? engineers (if the asset must be redesigned)?.
- what spares and tools are needed to do each task, including tools such as condition monitoring equipment.

It is only when resource requirements are clearly understood that we can decide exactly what systems are needed to manage the resources in such a way that the tasks get done, and hence that the functions of the assets are preserved.

This process can be likened to building a house. The foundations are the maintenance requirements of each asset, the walls are the resources required to fulfil the requirements (people/skills and spares/materials/tools) and the roof represents the systems needed to manage the resources (MMS).

To summarize, the development and execution of a maintenance strategy consists of three steps:
- formulate a maintenance strategy for each asset (work identification)
- acquire the resources needed to execute the strategy effectively (people, spares and tools)
- execute the strategy (acquire, deploy and operate the systems needed to manage the resources efficiently).

FIGURE 12

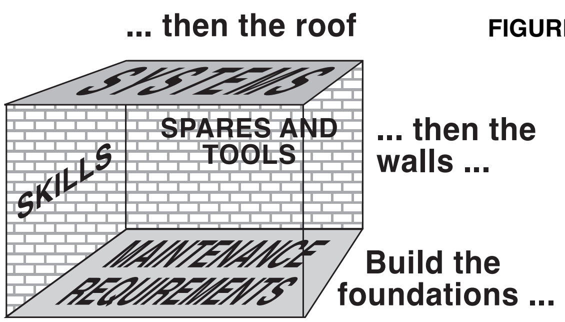

In other words, as shown in Figure 1, build your foundations first, then your walls, then your roof.

In terms of this structural analogy, it is worth noting that many enterprises spend immense amounts of time, energy and money on maintenance management systems (roofs) and on tools such as condition monitoring (part of the walls), but spend little or nothing on clarifying perceptions about what must really be done to cause the assets to continue to do what their users want them to do (the foundations). The result is elegant roofs and walls built over foundations that are the wrong shape, the wrong size, in the wrong place and not nearly strong enough to support the loads imposed upon them. The end result is a maintenance enterprise that is not nearly as effective as it should be.

This is not to suggest that we don't need an MMS or condition monitoring. Of course we do, in the same way that (nearly) every building needs a roof and walls. However, the roofs and walls must fit their foundations, and the foundation must be able to support the rest of the structure.

In essence, the only way to develop a truly viable, long-term maintenance strategy is to invest an appropriate amount of time and energy in *every* element of the process. In particular, avoid the temptation to concentrate too heavily or too soon on maintenance technologies and systems without first ensuring that everyone shares a clear, common and correct understanding about what must really be done to ensure that every asset continues to do what its users want it to do. In the absence of any comparable asset management strategy formulation processes, the only really effective way to do so is to arrange for groups of appropriately trained operators, maintainers, supervisors and specialists who live with the assets on a day-to-day basis to apply Reliability-centered Maintenance (RCM) under the guidance of a suitably qualified facilitator.

5 Reliability-centred Maintenance 5

Reliability-centered Maintenance, especially the derivative known as RCM2, incorporates all twelve of the new maxims discussed above. The following paragraphs briefly describe what this process consists of from a technical viewpoint, consider how it should be applied in a way that enables users to derive the maximum benefit quickly and economically, and provide a brief summary of what it achieves.

3.1 RCM: The seven basic questions

RCM is defined as 'a process used to determine *what must be done to ensure that any physical asset continues to do whatever its users want it to do* in its present operating context'. The RCM process entails asking seven questions about the asset or system under review, as follows:

- ***what are the functions and associated performance standards of the asset in its present operating context?***
- ***in what ways does it fail to fulfil its functions?***
- ***what causes each functional failure?***
- ***what happens when each failure occurs?***
- ***in what way does each failure matter?***
- ***what can be done to predict or prevent each failure?***
- ***what should be done if a suitable proactive task cannot be found?***

These questions are reviewed in the following paragraphs.

Functions and Performance Standards

Before it is possible to apply a process used to determine what must be done to ensure that any physical asset continues to do whatever its users want it to do in its present operating context, two things need to be done:

- determine what its users want it to do
- ensure that it is capable of doing what its users want.

This is why the first step in the RCM process is to define the functions of each asset in its operating context, together with the associated desired standards of performance. What users expect assets to be able to do can be split into two categories:

- *primary functions*, which summarize why the asset was acquired in the first place. This category of functions covers issues such as speed, output, carrying or storage capacity, product quality and customer service.
- *secondary functions*, which recognize that every asset is expected to do more than simply fulfil its primary functions. Users also have expectations in areas such as safety, control, containment, comfort, structural integrity, protection, operating efficiency, environmental integrity and even the appearance of the asset,

The users of the assets are usually in the best position to know exactly what contribution each asset makes to the physical and financial well-being of the organization, so they must be involved in the RCM process from the outset.

Functional Failures

The objectives of maintenance are defined by the functions and performance expectations of the asset under consideration. But how does maintenance achieve these objectives?

The only occurrence that is likely to stop any asset performing to the standard required by its users is some kind of failure. This suggests that maintenance achieves its objectives by adopting a suitable approach to the management of failure. However, before we can apply a suitable blend of failure management tools, we need to identify what failures can occur. The RCM process does this at two levels:

- by identifying what circumstances amount to a failed state
- then by asking what events can cause the asset to get into a failed state.

In the world of RCM, failed states are known as ***functional failures*** because they occur when an asset is *unable to fulfil a function to a standard of performance which is acceptable to the user*. In addition to the total inability to function, this definition encompasses partial failures, where the asset still functions but at an unacceptable level of performance (including situations where the asset cannot sustain acceptable levels of quality or accuracy). Clearly these can only be identified after the functions and performance standards of the asset have been defined.

Failure Modes

As mentioned in the previous paragraph, once each functional failure has been identified, the next step is to try to identify all the *events which are reasonably likely to cause each failed state*. These events are known as ***failure modes***.

'Reasonably likely' failure modes include those that have occurred on the same or similar equipment operating in the same context, failures which are currently being prevented by existing maintenance regimes, and failures that have not happened yet but that are considered to be real possibilities in the context in question.

Most traditional lists of failure modes incorporate failures caused by deterioration or normal wear and tear. However, the list should include failures caused by human errors (by operators or maintainers) and design flaws so that all reasonably likely causes of equipment failure can be identified and dealt with appropriately. It is also important to identify the cause of each failure in enough detail to make it possible to identify a suitable failure management policy.

Failure Effects

The fourth step in the RCM process entails listing ***failure effects***, which describe what happens when each failure mode occurs. These descriptions should include all the information needed to support the evaluation of the consequences of the failure, such as:

- what evidence (if any) that the failure has occurred
- in what ways (if any) it poses a threat to safety or the environment
- in what ways (if any) it affects production or operations
- what physical damage (if any) is caused by the failure
- what must be done to repair the failure.

The process of identifying functions, functional failures, failure modes and failure effects yields surprising and often very exciting opportunities for improving performance and safety, and for eliminating waste

Failure Consequences

Maxim 4 above pointed out that the main objective of proactive maintenance is not to avoid failure, but to eliminate, avoid or minimise the consequences of failures.

It is these consequences which most strongly influence the extent to which we try to prevent each failure. In other words, if a failure has serious consequences, we are likely to go to great lengths to try to avoid it. On the other hand, if it has little or no effect, then we may decide to do no routine maintenance beyond basic cleaning and lubrication.

A great strength of RCM is the extent to which it recognizes that the consequences of failures are far more important than their technical characteristics. The RCM process classifies these consequences into four groups, as follows:

- ***Hidden failure consequences:*** Hidden failures have no direct impact, but they expose the organization to multiple failures with serious, often catastrophic, consequences.
- ***Safety and environmental consequences:*** A failure has safety consequences if it could hurt or kill someone. It has environmental consequences if it could lead to a breach of any corporate, regional, national or international environmental standard.
- ***Operational consequences:*** A failure has operational consequences if it affects production (output, product quality, customer service or operating costs in addition to the direct cost of repair)

- ***Non-operational consequences:*** Evident failures that fall into this category affect neither safety nor production, so they involve only the direct cost of repair.

We see later how the RCM process uses these categories as the basis of a strategic framework for maintenance decision-making. By forcing a structured review of the consequences of each failure mode in terms of these categories, it integrates the operational, environmental and safety objectives of the maintenance function. This brings safety and the environment into the mainstream of maintenance management.

The consequence evaluation process shifts emphasis away from the idea that *all* failures are bad and must be prevented. It also encourages users to think more broadly about different ways of managing failure, rather than to concentrate only on failure prevention. Failure management techniques are divided into two categories:

- *proactive tasks:* these are tasks undertaken before a failure occurs, in order to prevent the item from getting into a failed state. They embrace what is traditionally known as 'predictive' and 'preventive' maintenance, although RCM uses the terms *scheduled restoration*, *scheduled discard* and *on-condition maintenance*
- *default actions:* these deal with the failed state, and are chosen when it is not possible to identify an effective proactive task. Default actions include *failure-finding, redesign* and *run-to-failure.*

Consequence evaluation is discussed again later. The following paragraphs look at proactive tasks in more detail

Proactive Tasks

Maxim 3 explained that many people still believe that the best way to optimize plant availability is to carry out overhauls or component replacements at fixed intervals, based on the assumption that most items operate reliably for a period 'X" then wear out. Maxim 3 went on to show that there is seldom a connection between reliability and operating age, and that unless the equipment suffers from a dominant age-related failure mode, age limits do little or nothing to improve the reliability of complex items.

An awareness of these facts has led some organizations to abandon the idea of proactive maintenance altogether. In fact, this can be the right thing to do for failures with minor consequences. But when the failure consequences are significant, *something* must be done to prevent or predict the failures, or at least to reduce the consequences. This brings us back to the question of proactive tasks. As mentioned earlier, proactive maintenance tasks can be divided in *preventive* maintenance and *predictive* maintenance.

Preventive tasks

Preventive tasks entail taking action at or before a specified age limit. This action can take one of two forms:

- scheduled restoration, which entails remanufacturing a component or overhauling an assembly at or before a specified age limit, regardless of its condition at the time
- scheduled discard, which entails discarding an item or component at or before a specified life limit, regardless of its condition at the time.

These used to be by far the most widely used forms of proactive maintenance. However for the reasons discussed above, they are much less widely used than they were twenty years ago (although they still have a place when there is hard evidence that they are appropriate.)

Predictive tasks

The continuing need to prevent certain types of failure, and the growing inability of classical preventive techniques to do so, are behind the growth of new types of failure management. The majority of these techniques rely on the fact that most failures give some warning of the fact that they are about to occur. These warnings are known as ***potential failures***, and are defined as *identifiable physical conditions which indicate that a functional failure is about to occur or is in the process of occurring.*

The new techniques are used to detect potential failures so that action can be taken to avoid the consequences which could occur if they degenerate into functional failures. They are known as *predictive maintenance, condition-based maintenance* or *on-condition maintenance* and include the array of techniques known as *condition monitoring*.

Used appropriately, predictive tasks are a very good way of managing failures, but they can also be an expensive waste of time. RCM enables decisions to be made with particular confidence about the technical feasibility of predictive maintenance tasks, and about whether they are worth doing. By applying the principles discussed under Maxim 9 above, RCM also enables appropriate decisions to be made about the frequency of these tasks.

Default Actions

RCM recognizes three major categories of default actions, as follows:

- *failure-finding:* As explained under Maxim 8, failure-finding tasks entail checking hidden functions periodically to determine whether they have failed.
- *redesign:* redesign entails making any one-off change to the built-in capability of a system. This includes modifications to the hardware and also covers once-off changes to procedures.
- *no scheduled maintenance:* as the name implies, this default entails making no effort to anticipate or prevent failure modes to which it is applied, and so those failures are simply allowed to occur and then corrected. This default is also called *run-to-failure.*

The RCM Task Selection Process

The RCM process applies a highly structured consequence evaluation and policy selection algorithm to each failure mode. This algorithm incorporates precise and easily understood criteria for deciding which (if any) of the proactive tasks is *technically feasible* in any context, and if so for deciding how often and by whom the tasks should be done. It also incorporates criteria for deciding whether any task is *worth doing,* a decision which is governed by how well the candidate task deals with the *consequences* of the failure. Finally, if a proactive task cannot be found that is both technically feasible and worth doing, the algorithm leads users to the most suitable default action.

This approach means that proactive tasks are only specified for failures that really need them, which in turn leads to substantial reductions in routine workloads. In fact, if RCM is correctly applied to existing maintenance programs, it reduces the amount of *routine* work (in other words, tasks to be undertaken on a *cyclic* basis) issued in each period, usually by 40% to 70%. On the other hand, if RCM is used to develop a new maintenance program, the resulting scheduled workload is much lower than if the program is developed by traditional methods. Less routine work also means that the remaining tasks are more likely to be done properly. This together with the elimination of counterproductive tasks leads to more effective maintenance.

3.2 Applying the RCM process

This section of this paper reviews key issues which must be managed when RCM is applied to any asset or system.

Planning

The successful application of RCM depends first and perhaps foremost on meticulous planning and preparation. The key elements of the planning process are as follows:

- Define the scope and boundaries of each project
- Define and wherever possible quantify the objectives of each project (now state and desired end state)
- Estimate the amount of time (number of meetings) needed to review the equipment in each area
- Identify project manager and facilitator(s)
- Identify participants (by title and by name)
- Plan training for participants and facilitators
- Plan date, time and location of each meeting
- Plan management audits of RCM recommendations
- Plan to implement the recommendations (maintenance tasks, design changes, changes to operating procedures)

Review groups

Maxims 10 and 11 stressed the need to involve shopfloor people and production/operations people in the maintenance strategy formulation process. This was because maintenance people and/or supervisors simply cannot answer all these questions on their own. Conversely, many (if not most) of the answers can only be supplied by production or operations people. This applies especially to questions about functions, desired performance, failure effects and failure consequences. For this reason, the RCM process should be applied by small teams which include *at least* one person from the maintenance function and one from the operations function.

The seniority of the group members is less important than the fact that they should have a thorough knowledge of the asset under review. Each group member should also have been trained in RCM. The make-up of a typical RCM review group is shown in Figure 13.

FIGURE 13

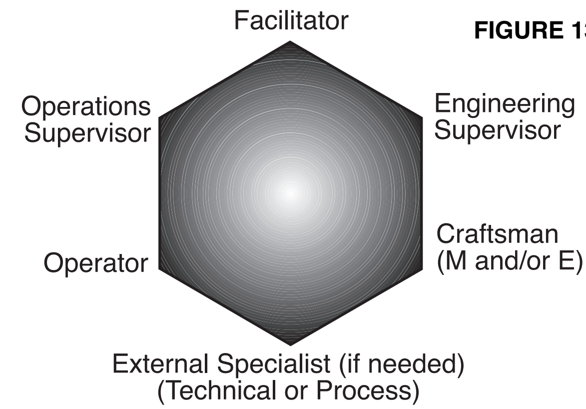

The use of these groups not only enables management to gain access to the knowledge and expertise of each member of the group on a systematic basis, but the members themselves gain a greatly enhanced understanding of the asset in its operating context.

Facilitators

RCM review groups work under the guidance of highly trained specialists in RCM, known as facilitators. The facilitators are the most important people in the RCM review process. Their role is to ensure that:

- the RCM analysis is carried out at the right level, that system boundaries are clearly defined, that no important items are overlooked and that the results of the analysis are properly recorded
- the questions are correctly understood by the group members, and that answers are recorded clearly and concisely
- the group reaches consensus in a brisk and orderly fashion, while retaining the enthusiasm and commitment of individual members
- the analysis progresses briskly and finishes on time.

Facilitators also work with RCM project managers or sponsors to ensure that each analysis is properly planned and receives appropriate managerial and logistic support.

The outcomes of an RCM analysis

If it is applied in the manner suggested above, an RCM analysis results in three tangible outcomes, as follows:

- schedules to be done by the maintenance department
- revised operating procedures for the operators of the asset
- a list of areas where one-off changes must be made to the design of the asset or the way in which it is operated to deal with situations where the asset cannot deliver the desired performance in its current configuration.

A less tangible but very valuable outcome is that participants in the process tend to start functioning much better as multi-disciplinary teams after their analyses have been completed.

Auditing
After the review has been completed for each asset, senior managers with overall responsibility for the equipment must satisfy themselves that the review is sensible and defensible. This entails deciding whether they agree with the definition of functions and performance standards, the identification of failure modes and the description of failure effects, the assessment of failure consequences and the selection of tasks.

Implementation
Once the RCM review has been audited and approved, the final step is to implement the tasks, procedures and one-time changes. The revised tasks and procedures must be drawn up in a way which ensures that they will be clearly understood and performed safely by the people to whom they are allocated. The maintenance tasks are then fed into suitable high- and low-frequency maintenance planning and control systems, while revised operating procedures are incorporated into standard operating procedure manuals. Modifications are usually dealt with by the engineering function.

3.3 What RCM Achieves

If RCM is correctly applied, it makes the following contributions to the performance of the organization:

- ***Greater safety*** and ***environmental integrity:*** RCM considers the safety and environmental implications of every failure mode before considering its effect on operations. This means that steps are taken to minimise all identifiable equipment-related safety and environmental hazards, if not eliminate them altogether. By integrating safety into the mainstream of maintenance decision-making, RCM also improves attitudes to safety.

- ***Improved operating performance (output, product quality and customer service):*** RCM recognizes that *all* types of maintenance have some value, and provides rules for deciding which is most suitable in every situation. By doing so, it helps ensure that only the most effective forms of maintenance are chosen for each asset, and that suitable action is taken in cases where maintenance cannot help. This much more tightly focused maintenance effort leads to quantum jumps in the performance of *existing assets* where these are sought.
  RCM was developed to help airlines draw up maintenance programs for new types of aircraft *before* they enter service. As a result, it is an ideal way to develop such programs for *new assets*, especially complex equipment for which no historical information is available. This saves much of the trial and error which is so often part of the development of new maintenance programs – trial that is time-consuming and frustrating, and error that can be very costly.

- ***Greater maintenance cost-effectiveness:*** By continually focusing attention on the maintenance activities that have most effect on the performance of the plant., RCM helps to ensure that everything spent on maintenance is spent where it will do the most good.

- ***Greater motivation of individuals,*** especially people who are involved in the review process. This leads to much wider 'ownership' of maintenance (and operations) problems and their solutions. It also means that solutions are more likely to endure.

- ***Better teamwork:*** RCM provides a common technical language for everyone who has anything to do with maintenance. This gives maintenance and operations people a better understanding of what maintenance can (and cannot) achieve and what must be done to achieve it.

- A ***comprehensive database:*** RCM ends with a comprehensive and fully documented record of the maintenance requirements of the assets reviewed. This makes it possible *to adapt to changing circumstances* (such as changing shift patterns or new technology) without having to reconsider all maintenance policies from scratch. It also enables equipment users to demonstrate that their maintenance programs are built on solid foundations (the *audit trail* required by more and more regulators). Finally, the information stored on RCM worksheets *reduces the effect of staff turnover* with its attendant loss of experience and expertise.
  An RCM review of the maintenance requirements of each asset also provides a much clearer view of the *skills required to maintain each asset*, and for deciding *what spares should be held in stock*. A valuable by-product is also *improved drawings and manuals.*

All of these issues are part of the mainstream of maintenance management, and many are already the target of improvement programs. A major feature of RCM is that it provides an effective step-by-step framework for tackling *all* of them at once, and for involving everyone who has anything to do with the equipment in the process.

RCM yields results very quickly. In fact, if they are correctly focused and correctly applied, RCM reviews can pay for themselves in a matter of months and sometimes even weeks The reviews transform both the perceived maintenance requirements of the assets used by the organization and the way in which the maintenance function as a whole is perceived. The result is more cost-effective, more harmonious and much more successful maintenance.

Part 3 of this paper proposed a maintenance mission statement which stressed that maintainers serve three distinct sets of customers: the owners of the assets, the users of the assets and society as a whole. Owners are satisfied if their assets generate a satisfactory return on investment. Users are satisfied if each asset conti-nues to do what they want it to do to standards of performance which they - the users - consider to be satisfactory. (In this context, satisfactory performance includes the notion that the risk of death or injury caused by equipment failure should be reduced to tolerable levels.) Finally, society is satisfied if the assets do not fail in ways that threaten the environment. Because they are maintaining assets on behalf of all these people, it could be said that maintainers are the custodians of the assets.

Double-entry Bookkeeping

In this context, parallels can be drawn between the custodianship of physical assets and the custodianship of financial assets. In 1494, a Florentine named Pacioli invented double-entry bookkeeping, the process at the heart of financial custodianship. To this day, throughout all branches of organized human endeavour, armies of bookkeepers and accountants use Pacioli's ideas to look after financial assets on behalf of the people who actually own, earn and spend the money. In their world, responsible custodianship means ensuring that all financial transactions are accounted for and the books balanced to the nearest penny at the end of every accounting period. The procedures and documentation needed to make this process work have become part of the way we are all obliged to do business, even though they are highly resource intensive and very expensive. Businesses the world over have learned that anything less precise quickly leads to financial chaos.

In the world of maintenance, our 'currency' is the failure mode. To exercise standards of custodianship similar to those of our financial brethren, we must ensure that every failure mode is properly 'accounted for'. This obliges us to exercise due diligence in trying to identify every failure mode which is reasonably likely to affect the functions of the assets in our care, to understand the consequences of each failure mode, to select the most cost-effective failure management policies, to deploy the most appropriate human and physical resources to execute the chosen policies and to ensure that each task is planned and executed in the right way, at the right time and by the right people.

The Consequences of Failure

Compare what happens when things go wrong in the worlds of financial and physical asset management. The worst consequences of the irresponsible custodianship of financial assets are that a business may go bankrupt and its custodians end up in prison. However, the worst consequence of the incorrect or irresponsible custodianship of physical assets is that people die, sometimes in very large numbers.

In fact, the extent to which the physical and financial health of most organizations now depends on their physical assets means that the pressure upon maintainers to exercise this custodianship in the most responsible fashion possible is becoming extraordinarily intense. Not only is this pressure arising from the expectations of the 'customers' of the maintenance service, but as explained in Maxim 1 above, it is attracting the increasingly demanding attention of regulatory bodies throughout the world.

In this environment, maintainers need to raise their standards of custodianship to far higher levels than have ever been acceptable in the past. And yet, at this point in time, industry in general still seems to be spending much more energy on the high precision management of its financial assets than of its physical assets, despite the fact that the consequences of incorrect custodianship are often far worse in the case of the latter than in the case of the former.

Experimentation

One reason why financial assets attract much greater senior management attention than physical assets is that the processes used to manage financial assets have been under development since Pacioli's era. By comparison, the concept of planned maintenance has been in existence for less than 50 years, while RCM was first codified barely 20 years ago. Terms like PdM (or CBM) and CMMS have only come into widespread use in the last 10 years. In short, industry is only just beginning to appreciate what must be done to exercise truly responsible custodianship of physical assets. We are decades away from establishing physical asset management processes that are as widely accepted and rigorously enforced as those used in the world of financial management.

Under these circumstances, it is not surprising that a great deal of experimentation is still going on in the world of physical asset management. Some of this experimentation is leading to developments that are of great value. In particular, think of the explosive growth in condition monitoring techniques, continuous advances in the CMMS field, rapidly growing understanding of the processes that cause systems to fail (including the part played by human error), and the formal incorporation of quantified risk into maintenance strategy formulation.

However, one area where we still have much to learn is in the field of RCM.

The History of RCM and the SAE RCM Standard

The RCM process finds its roots in work done by the international commercial aviation industry. Driven by the changes described in Part 2 of this paper, this industry needed to develop a comprehensive new process for deciding what maintenance work is needed to keep aircraft airborne. This process evolved steadily since its early beginnings in 1960. In its most recent form within aviation, it is known as MSG3 Revision 2.

In 1978, a report[1] was prepared for the US Department of Defense describing the then current state of this process. The report was written by F Stanley Nowlan and Howard Heap of United Airlines. It was entitled "Reliability-centered Maintenance", or RCM. In the early 1980's, RCM as described by Nowlan and Heap began to be used in industries other than aviation.

It soon became apparent that no other comparable technique exists for identifying the true, safe minimum of what must be done to preserve the functions of physical assets, especially in critical or hazardous situations. As a result, RCM has now been used by thousands of organizations spanning nearly every major field of organized human endeavour. It is becoming as fundamental to the practice of physical asset management as double-entry bookkeeping is to financial asset management.

The growing popularity of RCM has led to the development of numerous derivatives of this process. Some of these are refinements and enhancements of the process originally described by Nowlan and Heap. These include the processes used by the US Naval Air Command[2], the British Royal Navy[3] and RCM2[4].

However, less rigorous derivatives have also emerged, most of which are attempts to 'streamline' the maintenance strategy formulation process. Most of these attempts are made by well-intentioned people concentrating more on the cost of the strategy formulation process than on what it achieves. In fact, nearly all of the 'streamlined' maintenance strategy formulation processes encountered by the author to date contain logical or procedural flaws which increase risk to an extent that overwhelms any small advantage they may offer in reduced application costs. Chief among these processes are (1) those that attempt to combine two or all three of the incompatible methodologies needed to set intervals for different types of periodic maintenance tasks into one all-embracing formula, (2) those that place too much emphasis on assessing the 'criticality' of assets or systems before a detailed FMEA has been performed, and (3) those that reverse or simply skip key steps in the RCM process.

Ironically, it also transpires that many of these 'streamlined' techniques actually take longer and cost more to apply than the rigorous application of RCM, so even this small advantage is lost.

It is apparent to those who know RCM best that we all still need to learn much more about the intricate relationships between functions, failure mechanisms, failure consequences and failure management policies than we currently know. What is more, as mentioned earlier, the consequences of formulating inappropriate strategies are horrendous. It is a situation that demands more rigor, not less. As a result, too much emphasis on shortcuts right now is both dangerous and irresponsible.

A growing awareness of this issue led a group of the most concerned and experienced users of RCM from outside the world of commercial aviation to develop a set of criteria that can be used to assess whether any process that claims to be RCM is in fact RCM. Under the auspices of the International Society of Automotive Engineers (SAE) and spearheaded by the US Naval Air Command, this group produced a standard[5] published by the SAE in August 1999. It is entitled "Evaluation Criteria for Reliability-centered Maintenance (RCM) Processes".

This standard provides a yardstick that helps users to ensure that they are using a valid (and defensible) interpretation of the RCM process, and in so doing helps them to be sure that they are exercising truly responsible custodianship of their physical assets.

Auditing

A final point about responsible custodianship concerns auditing. In most organizations, financial managers have to submit their custodianship to exhaustive, expensive – and mandatory – external scrutiny at least once a year. At present, the notion of regular external audits of physical asset management activities is still in its infancy. However, the concept of an 'audit trail' is featuring in more and more industrial safety legislation. Our regulators are asking us not only to do the right things, but to be able to demonstrate in writing *why* we are doing them. The day is approaching when this will evolve into an audit process every bit as formalized and highly regulated as that to which our financial colleagues are subjected.

The depth, intrusiveness and cost of this audit process will be governed by the extent to which our regulators accept the validity of the methods we use to exercise custodianship of our physical assets, and the rigor and precision with which they consider us to be applying them. In short, if the people who own and use physical assets wish to maintain a reasonable degree of control over their own destiny, they must match if not exceed the standards of custodianship that are the norm in the world of financial asset management.

1 Nowlan FS and Heap H: *Reliability-centered Maintenance.* Springfield Virginia. National Technical Information Service, United States Department of Commerce
2 US Naval Air Systems Command: *NAVAIR 00-25-403: Guidelines for the Naval Aviation Reliability Centered Maintenance Process.* Philadelphia, Pennsylvania, USA. US Department of Defense Publications
3 RCM Implementation Team, Royal Navy: *NES 45 Naval Engineering Standard 45, Requirements for the Application of Reliability-Centred Maintenance Techniques to HM Ships, Royal Fleet Auxiliaries and other Naval Auxiliary Vessels (Restricted-Commercial)*: Foxhill, Bath, United Kingdom. UK Ministry of Defence Publications
4 Moubray JM: *Reliability-centered Maintenance*: New York, New York USA: Industrial Press
5 International Society of Automotive Engineers: *Standard JA1011 - Evaluation criteria for Reliability-Centered Maintenance (RCM) processes:* Warrendale, Pennsylvania, USA: SAE Publications