

THE RESPONSIBLE CUSTODIANSHIP OF PHYSICAL ASSETS

Presented by
ALADON Ltd
in collaboration with
John Moubray

1 A Maintenance Mission

What exactly is the purpose of the maintenance function? In a world of growing expectations, increasingly onerous regulatory constraints, shifting technological paradigms and endless reorganisations – all of which must be dealt with urgently – it is easy to get lost. In this environment, just as most major corporations develop formal mission statements to help them maintain a steady course through an ocean of distractions, it is worth developing a formal mission statement to help maintenance do likewise.

Perhaps a good place to start would be to look at the meaning of the word ‘maintain’. The Oxford dictionary defines *maintain* as *cause to continue*.

Cause what, you may ask, to continue to do what? The first ‘what’ is easy. Maintenance exists because we have *physical assets* which need maintaining. So the mission statement must reflect the fact that maintenance is first and foremost about physical assets.

But what is it that they must continue to do? The answer lies in the fact that every physical asset is put into service because someone wants it to do something. In other words, it is expected to fulfil a specific function or functions. So it follows 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 and preventive maintenance, failure-finding, 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 organisations. 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 stress the need to make the most cost-effective choices from the full array of options.

When considering failure management options, note that failures only attract attention because they have consequences. Failures can affect output, safety, environmental integrity, output, product quality, customer service, protection and operating costs in addition to repair costs. The severity and frequency with which a failure incurs these consequences 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 minimised throughout their useful lives, not just to the end of the next accounting period.

Finally, the mission statement must recognise 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.***

2 Developing A Maintenance Strategy

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 do we start? Buy a new maintenance management system (MMS)? Reorganise? Invest in loads of condition monitoring equipment? Knock the whole place down and rebuild it?

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.

At this point, we have decided what must be done to preserve the functions of our assets. This process could be called “work identification”.

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 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 correctly, 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 needed to fulfil the requirements (skills and spares/tools) and the roof represents the systems needed to manage the resources (MMS).

Looking at maintenance requirements in the context of the *functions* of each asset (by seeking to preserve what the asset *does* rather than what it *is*), completely transforms the way in which the requirements are perceived. In other words, such a review changes the size, shape and location of the foundations upon which the maintenance enterprise is built. Clearly, when the foundations change, everything built on those foundations must also change.

The good news is that if the review of requirements - the work identification process - is carried out correctly, the foundations not only end up somewhere else, but they are usually much smaller than if requirements are determined by old fashioned seat-of-the-pants methods.

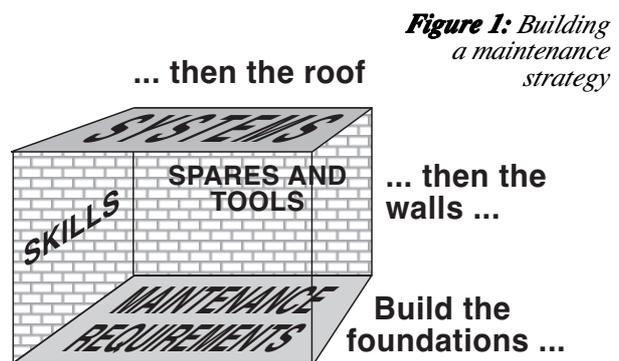
Smaller foundations mean that the entire structure (resources and systems) built on those foundations will also be smaller.

Even better news is that the initial focus on functions makes the whole enterprise far, far more effective.

To summarise, 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).

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



3 Building Strong Foundations

As every builder knows, the integrity of any structure depends first and foremost on the integrity of its foundations. So if we seek a maintenance enterprise that is robust enough to satisfy all the expectations of its customers, then:

- its foundations must always be the right size and shape, and in the right place
- the foundations must be sufficiently solid to bear all the loads placed upon them.

Building solid foundations means that the building project must be planned properly, the ground must be prepared correctly, the foundations must be properly designed, the right materials used and the foundations must be built by people with appropriate knowledge and skills.

Planning the project means that clear objectives must be established, resources allocated and a plan prepared.

Preparing the ground means that everyone in the organisation served by the maintenance enterprise must clearly understand what maintenance can and cannot achieve, and what they must do to help to achieve it. *Designing the foundations* and *selecting the right materials* means systematically defining the functions and required performance standards of each asset, deciding what failure modes are reasonably likely to cause it to fail, assessing the effects and consequences of each failure, and selecting a failure management policy that deals appropriately with the consequences.

Using *appropriate people* means that the exercise must be performed by groups of people who have a thorough understanding of each asset in its operating context, working under the guidance of someone who profoundly understands the process being used to assess the maintenance requirements and who has a long-term vested interest in the success of the project.

In the absence of any comparable asset management strategy formulation processes, the only really effective way to do all this at once for modern, complex industrial processes is to arrange for groups of appropriately trained operators, maintainers, supervisors and specialists who live with the asset on a day-to-day basis to apply Reliability-centred Maintenance (RCM) under the guidance of a suitably qualified facilitator.

4 Reliability-centred Maintenance

Reliability centred Maintenance 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’. It entails asking seven questions about the asset 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 if a suitable proactive task cannot be found?***

These questions are reviewed in the following paragraphs.

4.1 Functions and Performance Standards

Part 2 of this paper mentioned that it is only when the functions of an asset have been defined that it becomes clear exactly what maintenance is trying to achieve, and also precisely what is meant by “failed”.

For this reason, 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. The users of the assets are usually in by far the best position to know exactly what contribution each asset makes to the physical and financial well-being of the organisation as a whole, so it is essential that they are involved in the RCM process from the outset.

4.2 Functional Failures

The objectives of maintenance are defined by the functions and associated performance expectations of the asset. 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. 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:

- firstly, 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).

4.3 Failure Modes

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 that are currently being prevented by existing maintenance tasks, 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 (on the part of operators and 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 for it to be possible to identify a suitable failure management policy.

4.4 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 failure consequences, 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.

4.5 Failure Consequences

A detailed analysis of an average industrial undertaking is likely to yield between three and ten thousand possible failure modes. As mentioned in Part 1 of this paper, each of these failures affects the organisation in some way, but in each case, the consequences are different. The RCM process classifies failure consequences into four groups, as follows:

- **Hidden failure consequences:** Hidden failures have no direct impact, but they expose the organisation to multiple failures with serious consequences.
- **Safety and environmental consequences:** A failure has safety consequences if it could hurt or kill someone. It has environmental consequences if it could breach a 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 operations, so they involve only the direct cost of repair.

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 the above categories, it focuses attention on the maintenance activities which have most effect on the performance of the organisation, and diverts energy away from those that have little or no effect (or which may even be actively counterproductive). It also encourages users to think more broadly about different ways of managing failure, rather than to concentrate only on failure prevention.

4.6 Failure Management Policy Selection

Failure management policies fall 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. As discussed below, RCM further subdivides these tasks into *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*.

Scheduled restoration and scheduled discard tasks

Scheduled restoration entails remanufacturing a component or overhauling an assembly at or before a specified age limit, regardless of its condition at the time. Similarly, scheduled discard entails discarding an item at or before a specified life limit, regardless of its condition at the time. Collectively, these two types of tasks are now generally known as *preventive maintenance*.

On-condition tasks

On-condition 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 that indicate that a functional failure is about to occur or is in the process of occurring*.

On-condition tasks are used to detect potential failures so that action can be taken to reduce or eliminate the consequences that could occur if they were to degenerate into functional failures. This category of tasks includes all types of *predictive maintenance*, *condition-based maintenance* and *condition monitoring*.

Failure-finding

Failure-finding entails checking hidden functions to find out whether they have *failed* (as opposed to on-condition task, which entail checking if something is *failing*).

Redesign

Redesign entails making any one-time change to the built-in capability of a system. This includes changes to hardware, one-time changes to procedures and if necessary, training.

No scheduled maintenance

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, then repaired. This default is also called *run-to-failure*.

4.7 The RCM Task Selection Process

The RCM process applies a highly structured consequence evaluation and policy selection algorithm to each failure mode. It 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 that 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 for dealing with the failure.

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 done 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.

4.8 Applying RCM

Correctly applied, RCM contributes to remarkable improvements in maintenance effectiveness, and often does so surprisingly quickly. However, as with any fundamental change management project, RCM only succeeds if proper attention is paid to thorough planning, how and by whom the analysis is performed, auditing and implementation. These issues are discussed in the following paragraphs

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

We have seen that the RCM process embodies seven basic questions. In practice, maintenance people simply cannot answer all these questions on their own. This is because many (if not most) of the answers can only be supplied by production or operations people. This applies especially to questions concerning functions, desired performance, failure effects and failure consequences.

For this reason, a review of the maintenance requirements of any asset should be done 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 2.

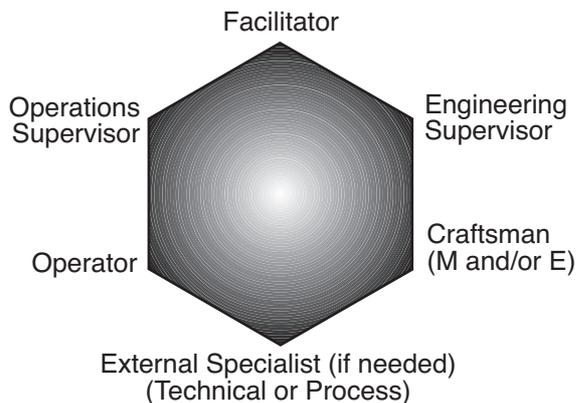


Figure 2:

A typical RCM review group

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 learn a great deal about how the asset works.

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
- RCM is correctly understood and applied by the group
- the group reaches consensus in a brisk and orderly fashion, while retaining their enthusiasm and commitment
- the analysis progresses as planned 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 asset operators
- 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 are 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.

4.9 What RCM achieves

The most important single contribution of the RCM process to industry is that if it is correctly applied, it provides a far more solid foundation for the maintenance enterprise than anything which has been available hitherto. Key areas in which it contributes directly to maintenance effectiveness and efficiency are as follows:

- **Greater safety and environmental integrity:** RCM considers the safety and environmental implications of every failure mode before considering its effect on operations. This brings safety and the environment into the mainstream of maintenance decision-making
- **Improved operating performance (output, product quality, customer service):** By concentrating on what physical assets *do* (their functions) rather than what they *are*, RCM enables users to identify much more clearly and precisely what must be done to achieve real and substantial long-term improvements in plant availability and reliability
- **Greater maintenance cost-effectiveness:** RCM continually focuses attention on the maintenance activities which have most effect on the performance of the plant. This helps to ensure that everything spent on maintenance is spent where it will do the most good.
- **Longer useful life of expensive items,** due to a careful focus on the use of on-condition maintenance.
- **A comprehensive database:** An RCM review ends with a comprehensive and fully documented record of the maintenance requirements of all the significant assets used by the organisation. This makes it possible *to adapt to changing circumstances* without having to reconsider all maintenance policies from scratch. It also enables equipment users to demonstrate that their maintenance programs are built on rational foundations (the *audit trail* required by more and more regulators).
- **Greater motivation of individuals,** especially people who are involved in the review process. This leads to greatly improved general understanding of the equipment in its operating context, together with wider 'ownership' of maintenance problems and their solutions. It also means that solutions are more likely to endure.
- **Better teamwork:** RCM provides a common, easily understood 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.

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..

In terms of our structural analogy, it is worth noting that many maintenance 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 appropriate amounts of time and energy in *every* element of the process. In particular, avoid the temptation to concentrate too soon or too heavily on maintenance techniques and systems without first ensuring that everyone shares a clear, common and correct understanding of what must be done to ensure that every asset continues to do what its users want it to do.

5 Responsible Custodianship

The first part of this paper proposed a maintenance mission statement. In doing so, it stressed that 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 investment. Users are satisfied if each asset continues 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 which 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.

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 organised 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 that is reasonably likely to affect the functions of our assets, 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.

In the context of this analogy, 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 organisations now depends on the continued physical and functional integrity of their 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 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 they apply if we are thought to have got it wrong are also becoming steadily more ferocious. For example, the British government has recently introduced a new class of crime called 'corporate manslaughter', to be applied to the senior executives of organisations where fatalities can be shown to be the result of irresponsible custodianship.

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 spends 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.

This is partly because 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 Reliability-centred Maintenance was first codified in the Nowlan & Heap report 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 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 which 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.

One area where we still have a great deal to learn is in the field of RCM. It has been extensively applied with great success by the aviation industry in particular, in addition to which the author and his associates have been involved in the application of RCM to physical assets on more than 1000 sites encompassing nearly every major field of organised human endeavour. As a result, the process is well established. However, despite the tremendous successes enjoyed by those who apply RCM correctly, industry in general is only just starting to come to grips with it.

One feature of this learning process is the number of attempts that are being made to 'streamline' the maintenance strategy formulation process. Most of these attempts are being made by well-intentioned people concentrating more on the cost of the strategy formulation process than on what it achieves. However, 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, so too much emphasis on shortcuts right now is both dangerous and irresponsible.

(In fact, nearly all of the 'streamlined' maintenance strategy formulation processes encountered by the author to date contain logical or procedural flaws that 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 the three 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. So if we wish to be truly responsible custodians of our physical assets, we need to recognise that shortcuts simply have no place in the maintenance strategy formulation process in general, and in the application of RCM in particular.)

A further point about responsible custodianship concerns auditing. In most organisations, 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 formalised 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 how much our regulators accept the validity of the methods we use to exercise custodianship of our physical assets, and the rigour and precision with which they consider us to be applying them. In short, if the world of physical asset management wishes to maintain a reasonable degree of control over its own destiny, it must match if not exceed the standards of custodianship that are the norm in the world of financial asset management. Right now, how many of us can honestly say that it does?

About John Moubray:

John Moubray is a graduate mechanical engineer. He lives in Asheville, North Carolina, and is Chief Executive of Aladon Ltd. After graduating in 1971, he worked for two years as a maintenance planner in a packaging plant and for one year as a commercial field engineer for a major oil company.

In 1974, he joined a large multi-disciplinary management consulting company. He worked for this company for twelve years, specialising in the development and implementation of manual and computerised maintenance management systems for a wide variety of clients in the mining, manufacturing and electric utility sectors.

He first encountered Reliability-Centred Maintenance in 1981 and soon became aware of its universality and the extent to which it fills a void in general industrial maintenance thinking. This led him to seek out Stan Nowlan – co-author of the 1978 Nowlan & Heap report on RCM – in 1982. Under the guidance and mentorship of Stan Nowlan (which continued until Stan's death in 1995), Moubray began to develop a suite of training and

support services designed to transfer the technology of RCM to industrial clients.

In 1986, Moubray decided to concentrate exclusively on RCM. This led him to found Aladon Ltd in 1986. Aladon is a UK-based consulting company specialising in the application of RCM. Together with a world-wide network of licensees, Aladon has helped clients to apply RCM on more than 1000 sites in 40 countries.

The first edition of Moubray's book entitled "Reliability-centred Maintenance" was published in 1991 in the UK and in 1992 in the USA. The second edition was published in both countries in 1997, and more than 40 000 copies of this book are now in circulation in several languages. He has published numerous articles on RCM and has addressed conferences in Europe, the Americas, the Western Pacific and Africa.

His article entitled "Maintenance Management - A New Paradigm" has been published in leading maintenance journals in 16 countries.