

Captured by Data

A paper by Daryl Mather - Specialist in physical asset management, author of the book, *The Maintenance Scorecard*, ISBN 0831131810 darylm@strategic-advantages.com

Enterprise Asset Management Systems (EAM) and the aims of modern maintenance

Since the late 1980's EAM vendors throughout the world have pitched their products based partly on the ability to capture, manipulate, and analyze, historical failure data. Part of the stated benefits case is often the ability to highlight the causes for poor performing assets, provide the volume and quality of information for determining how best to manage the assets, and informing decisions regarding end-of-life and other investment points.

This benefits case covers the principal drivers for most maintenance managers today and it has been used to justify millions of dollars worth of investment, and has placed the modern EAM system at the centre of corporations that are driving to improve asset performance. On the surface it appears to be a logical approach for problems relating to asset performance, and using this approach companies do, of course, achieve results.

The implementation of these products, when bought for these reasons, often focuses on optimizing processes to capture the dynamic data on asset failures, which is then used throughout the system. MRO¹ style inventory management algorithms, for example, use this information as one of the key inputs to determine minimum stocking levels, reorder points and the corresponding reorder quantities.

If we want to understand the validity of this line of thinking it is necessary to first explore the aims of maintenance, and how asset data can be used to further those aims.

“Maintenance” is a term generally used to define the routine activities to sustain standards of performance throughout the in-service, or operational, part of the asset life cycle. In doing this, the maintenance policy designer needs to take account of a range of factors. These include the complexities of operating environment, the available resources for performing maintenance, and the ability of the asset to meet its current performance standards.

In the past, this would be the extent of the maintenance analysts' role. One of the realities they face is that at times assets are under a demand greater than, or extremely close to, their inherent capabilities. As a result analysts often find themselves recommending and analyzing activities of not only maintenance, but also other areas of **asset management**, namely those of asset modification and operations.

Safety and environmental compliance play their part in creating the drive for this activity, particularly given the changing legal and regulatory frameworks around these two areas; in some industries they are even the principal drivers. However for most businesses the goal remains that of maximum value from their investment. This means getting the maximum performance possible from the assets, for the least amount spent.

In the original report and appendices that produced Reliability-centered Maintenance (RCM) the authors defined critical failures, initially, as those failures with an impact on safety. Today the term “critical failure” is often used to group failures that will cause what companies consider to be high-impact consequences, a definition that is too variable for a general discussion. For the sake of simplicity, “critical failure” in this paper refers to **all** failures that will cause the asset to perform to a standard less than what is required of it.ⁱⁱ


If an asset management program is aimed at maximum cost-effectivenessⁱⁱⁱ over an assets life, then it **must** look at the management of critical failures. By definition, this approach is centered on the reliability of the asset. (Or reliability-centered)

So, in essence, the role of the policy designer can be defined as the formulating cost-effective asset management programs, routine activities and one-of procedural and design changes, to maintain standards of performance through reducing the likelihood of critical failures to an acceptable level, or eliminating them. This is also the essence of modern RCM.

The data dilemma

Immediately we start to see a contradiction between the aims of maintenance, and those often quoted of EAM systems. Non-critical failures are those of low or negligible cost consequences only. These are **acceptable** and can be allowed to occur. Therefore a policy that focuses on data capture and later analysis as its base can be used effectively. Over time the level of information will accumulate to allow asset owners, and policy designers, to determine the correct maintenance policy with a high degree of confidence.

Acceptable and unacceptable failures



However, critical failures, those that cause an asset to under perform, have **unacceptable** consequences and cannot always be managed in a similar way. For example, if a failure has high operational impact or economic consequences, then allowing it to fail prior to determining how to manage them is actively counterproductive to the aims of cost effective asset management. Moreover, recent history reinforces the fact that failure of assets can lead to consequences in safety^{iv} or breaches of environmental regulations.^v

So, if our policy for determining how best to manage physical assets is based around data capture, then we are creating an environment that runs counter to the principles of responsible asset stewardship in the 21st century.

The underlying theories of maintenance and that of reliability are based on the theory of probability and on the properties of distribution functions that have been found to occur frequently, and play a role in the prediction of survival characteristics.^{vi}

Critical failures are, by their very nature, serious. When they occur they are often designed out, a replacement asset is installed, or some other initiative is put in place to ensure that they don't recur. As a result the volume of data available for analysis is often small, therefore the ability of statistical analysis to deliver results within a high level of confidence is questionable at best.

This fundamental fact of managing physical assets highlights two flaws with the case of capturing data for designing maintenance programs. First, collecting failure information for future decisions means managing the asset base in a way that runs counter to basic aims of modern maintenance management. Second, even if a company was to progress down this path, the nature of critical failures is such that they would not lend themselves to extensive statistical review.

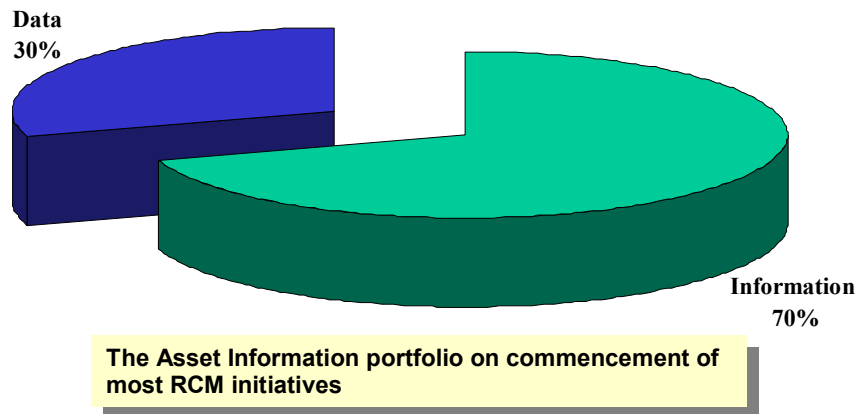
By establishing an effective, or reliability centered, maintenance regime, the policy designer is in effect creating a management environment that attempts to **reduce** failure information, **not** increase it. The more effective a maintenance program is, the fewer critical failures will occur, and correspondingly less information will be available to the maintenance policy designer regarding operational failures.^{vii} The more optimal a maintenance program is, the lower the volume of data there will be.

Designing maintenance policy

When maintenance policy designers begin to develop a management program they are almost always confronted with a lack of reliable data to base their judgments on. It has been the experience of the author that most companies start reliability initiatives using an information base that is made up of approximately 30% hard data, and 70% of knowledge and experience.

One of the leading reasons for this is the nature of critical failures and the response they provoke. However there are often other factors such as data capturing processes, consistency of the data, and the tendency to focus efforts in areas that are of little value to the design of maintenance policy. With EAM technologies changing continually, there are often upgrade projects, changeover projects, and other ways that data can become diluted.

Figure 1. Corporate Knowledge = Data + Information



There are still other key reasons why data from many EAM implementations are of limited value only. Principal among these is the fact that even with well-controlled and precise business processes for capturing data, some of the critical failures that will need to be managed may not yet have occurred. An EAM system, managing a maintenance program that is either reactive or unstructured, will only have a small impact on a policy development initiative.

At best they may have collected information to tell us that faults have occurred, at a heavy cost to the organization, but with small volumes of critical failures and limited information regarding the causes of failure. RCM facilitates the creation of maintenance programs by analyzing the four fundamental causes of critical failures of assets, namely:

- ❑ Poor asset selection (Never fit for purpose)
- ❑ Asset degradation over time (Becomes unfit for purpose)
- ❑ Poor asset operation (Operated outside of the original purpose)
- ❑ And, exceptional human errors (Generally following the GEM^{viii} principles)

The RCM Analyst needs to analyze all of the reasonably likely^{ix} failure modes in these four areas, to an adequate level of detail. Determining the potential causes for failures in these areas, for a given operating environment, is in part informed by data, but the vast majority of the information will come from other sources.

Sources such as operators' logs are strong sources for potential signs of failure, as well as for failures often not found in the corporate EAM. Equipment manufacturers' guides are also powerful sources for gleaning information regarding failure causes and failure rates. However, all information from a manufacturer needs to be understood in the context of how you are using the asset, and the, often conservative, estimates of the manufacturer. For example, if there are operational reasons why your pumping system is subject to random foreign objects, for whatever reason, then failure rates for impeller wear can become skewed.

Other sources of empirical data can be found in operational systems such as SCADA or CITECT, commercial databanks, user groups, and at times consultant organizations. Similarly to information from manufacturers there is a need to understand how this applies to the operating environment of your assets. As asset owners require more and more technologically advanced products, items come onto the market with limited test data in operational installations, further complicating the issues of maintenance design through data.

The factors that decide the lengths that an RCM Analyst should go to collect empirical data is driven by a combination of the perceived risk, (probability x consequence), and of course the limitations set on maintenance policy design by commercial pressures. Even when all barriers are removed from the path of the RCM Analyst, they are often confronted faced with an absence of real operational data on critical failures.

The vast majority of the information regarding how the assets are managed, how they can fail, and how they should be managed, will come from the people who manage the assets on a day-to-day basis. Potential and historic failure modes, rate of failure, actual maintenance performed (not what the system says, but what is **really** done), why a certain task was put into place in the first place, and the operational practices and reasons why, are all elements of information that are not easily found in data, but in knowledge.

This is one of the overlooked side-benefits of applying the RCM process, that of capturing **knowledge**, not merely data. As the workforce continues to age, entry rates continue to fall in favor of other managerial areas, and as the workforce becomes more mobile, the RCM process, and the skills of trained RCM Analysts, provides a structured method to reduce the impact of diminishing experience.

RCM and the role of the EAM

Among the areas where modern EAM systems **do** provide substantial benefits is through driving out inefficiencies in business processes. Through the capture, storage, manipulation, and display of historical transactional data, companies can take great leaps forward in the level of efficiency with which they execute maintenance programs. For example, through ensuring that delays in executing work are captured analyzed and resolved, or by being able to display trends in performance and cost over time.

The effectiveness of a maintenance task comes from how it manages failure modes, not from the level of efficiency that it is executed with. The original RCM study revealed that many routine tasks could actually contribute to failure, or to lower cost-effectiveness by having limited or no impact on the performance of the asset. (In effect wasting the maintenance budget) Executing these tasks with greater efficiency would have either have no impact at all on effectiveness, or possibly even magnify the effects of unsuitable tasks.

For example, after a lot of time working with a utility company in the UK it became clear that the reported schedule compliance was not an accurate figure. Schedules were regularly coming in with 100% compliance, while the reality was that they were actually performing at around 25%.


After some investigation it turned out that the crafts people recognized the most of the regimes that were coming out of the system were either counter productive, or not

applicable at all. So they were fortunately omitted. Prior to installing the EAM system they were working with job cards in separate systems, once the EAM went “live” these were collated and assigned to all similar assets regardless of operational context.

This is where RCM style methodologies contribute to the modern EAM or CMMS system. By providing the content that the system needs to manage, they are ensuring that the right job is being executed in the right way. This is common sense, and practitioners of RCM have been emphasizing this point for many years.

What is often not emphasized, however, is that having an effective maintenance program in place, and integrated with the EAM system ensures that future efforts of data capture are executed in a manner that supports the principles of responsible asset stewardship. The effect of building a data capture program on the back of an effective maintenance program is to **reverse**, over time, the ratio of hard-data to human-knowledge that is available for decision making.

Figure 2. Integration of EAM and Reliability-centered Maintenance



On performing the analysis, the structured approach within the decision diagram drives RCM Analysts to develop an asset management program that is practical, cost effective, and tailored to a given level of performance and risk. There are two main outputs of any correctly performed analysis. The first is one-of changes to; procedures, software, asset configurations, asset types, company policies and asset designs.

The second area is a group of routine maintenance tasks designed to manage the failure mode under analysis. Aside from combinations of policies, RCM supports 5 different maintenance policy options^x as detailed below. These make up the bulk of the content that the EAM system is installed to manage.

1. Predictive Maintenance (**P_{Tive}**) – a task to **predict** when a failure mode is about to occur.
2. Preventive Restoration (**P_{Res}**) – a task to **prevent** failure through applying a task, at a time or usage based interval, to restore the assets’ original resistance to failure.
3. Preventive Replacement (**P_{Rep}**) – a task to **prevent** failure through replacing an asset or component, at a time or usage based interval.
4. Detective Maintenance – (**D_{Tive}**) a task to **detect** whether an item has failed or not. This task is only applied to failure modes that RCM classifies as hidden.
5. Run-to-Fail – (**RTF**) a policy to allow an asset to **fail**, rather than applying any form of routine maintenance. Failure modes that are allowed to run-to-failure have low, or negligible, consequences in terms of cost only. These are non-critical, or acceptable, failures that were referred to earlier in this document.

An RCM based process selects these tasks based on their **applicability** and **effectiveness** as defined within the decision algorithms. These issues have been commented on many times and will not be dealt with in great detail within this paper.

For modern RCM Analysts the routine maintenance tasks are of interest not only because of the impact they have on asset performance, but also because the way they can be used to develop the asset information portfolio, contribute to whole-of-life costing, and to provide an additional tool for proactive monitoring of asset performance and corporate risk exposure.


As with the logic of the decision diagram, the criteria and characteristics of each of these policy choices has been detailed many times, and it is not necessary to describe them in detail in this paper. However, it is necessary to detail how they affect the collection, management, and use of dynamic asset data.

Predictive Maintenance

As detailed in **Figure 3** below, Predictive Maintenance (P_{Tive}) tasks are established to try to detect the warning signs that indicate the onset of failure, thus allowing for actions to be taken to avoid the failure. Yet there is also another aspect of P_{Tive} tasks that is often overlooked. That of the corrective, or **Predicted** (P_{Ted}), task once warning signs have been detected.

Immediately following the analysis, the information established at this point can be used for creating **proactive** whole-of-life costing models that are directly tied to performance and risk.

Figure 3. Tasks involved in predictive maintenance



Whole-of-life cost of an asset, or component, subject to Predictive Maintenance tasks

$$= (\text{Cost } (P_{Tive}) \times n) + \text{Cost } (P_{Ted})$$


Where n represents the number of times the P_{Tive} task is likely to be executed. This also drives estimates of the time between installation and likely failure. It needs to be recognized that the corrective, or P_{Ted} , task is executed at a time less than end-of life. (Although small)

As time passes the amount of data that is collected on these tasks will grow, collected now in a responsible manner, and can also be used in statistical models regarding asset degradation and predictions of capital spend requirements. By the inclusion of these outputs of an RCM analysis, asset managers can use the results with increasing confidence as predictors of whole of life cost profiles, and end of life points.

Preventive Maintenance

Where Predictive Maintenance tasks cannot be applied, for whatever reason, the next two options on either side of the decision diagram are Preventive Maintenance tasks. These are tasks that are aimed at either restoring an assets resistance to failure (P_{Res}), or replacing the asset at a time before the failures can occur. (P_{Rep}) Thus preventing failures. These tasks have limited use and are based on age, usage, or some other representation of time.

Figure 4. Tasks involved in preventive maintenance



When applied correctly these tasks are part of the approach to maintenance that, by necessity, reduces the volume of failure data available for statistical analysis. However, with the component out of the operational environment, it can safely be tested to try to establish the extent of its remaining economically useful life.

Whole-of-life cost of an asset, or component, subject to Preventive Maintenance tasks

$$(\text{Cost } (P_{Res}) - \text{or} - (\text{Cost } (P_{Rep}))$$

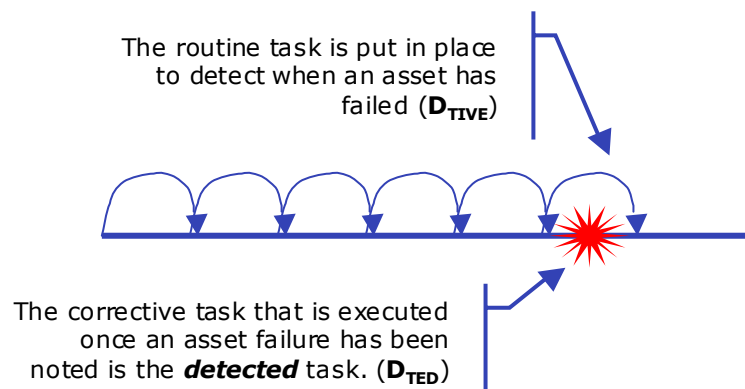
This is an additional task and one that would not be generated from the RCM analysis. Yet it represents another aspect of responsible data capture and is an important element of businesses where confidence in statistical life prediction, and whole of life costing models, are of importance.^{x1}

Detective Maintenance

As with predictive maintenance tasks there are actually two tasks that are being implemented here. First the **detective** (D_{Tive}) maintenance task, and second the **detected** (D_{Ted}) maintenance task. The result of this is the same as with the predictive maintenance tasks. That is, it provides further information about the likely failure rate,

collected in a responsible manner, which can be used to inform decisions regarding optimizing the frequency of this task.

Figure 5. Tasks involved in detective maintenance



Whole-of-life cost of an asset, or component, subject to Detective Maintenance tasks

$$= (\text{Cost } (D_{TIVE}) \times n) + \text{Cost } (D_{TED})$$

Where n represents the number of times the D_{TIVE} task is likely to be executed. This also drives estimates of the time between installation and likely failure. It needs to be recognized that the corrective, or D_{TED} , task is executed at a time **greater** than end-of life due to the characteristics of this task. As time passes the data collected can be used to inform decisions and whole-of-life models with increasing certainty.

This is particularly relevant for hidden failures, or hidden functions as they are sometimes called. When implementing the outcomes of an RCM analysis, some of the tasks are D_{TIVE} tasks. That is, they are tasks put in place to detect if a failure has occurred. Often, the items being tested have not been tested for a long period of time, sometimes years. And often nobody knows if they are working or not!


So when establishing the initial D_{TIVE} task frequencies, often the information used is not very certain and backed by only the experiences and memories of those involved in the exercise. Fortunately manufacturers **do** often have a good level of information regarding failure rates in these sorts of devices. But the result is still quite conservative and not tuned for the specific operational climate. Performing D_{TIVE} tasks will immediately help the company to establish some baseline information regarding failure rates of the device.

Run-to-Failure

The last policy option, aside from redesign and combinations of tasks, is that of run-to-failure. This option is for the acceptable, or low / negligible cost, failures detailed in figure 1. The EAM will allow these failures to be captured for analysis to inform whole of life cost models, spending forecasts, and to be used in reviewing maintenance policies when relevant.

Figure 6. Tasks involved in Run-to-Failure policies

The only task in a run-to-fail policy is the corrective task that is executed once the failure occurs.



Along with the responsible data capture forced by these policy options, configuring and managing the EAM in line with RCM thinking will also allow visibility of *exceptional failures*.

Due to the way that RCM is, by necessity, carried out, there is the possibility that some failures may be missed. Modern methods of execution have expanded the original default method of team-based analyses to include expert analysis sources outside of the team, but there always remains the possibility that the analysis will miss a critical failure despite the best efforts of the analyst and those involved.

In these circumstances the data recorded in these **exceptional** failures will provide the impetus for the analyst to revisit the analysis to factor in this failure mode and to put in place a relevant management policy. It is not an area that is used for capturing data for statistical analysis and is, as the name suggests, the exception rather than the rule.

It can be seen that part of the role of the modern RCM Analyst is not only to *minimize* the volume of failure data that is collected for later analysis, but also to *maximize* the quality and usability of data that is captured via collection methods that support the principles of responsible asset stewardship. It can also be seen that advances in modern technology, combined with the growing needs of asset intensive companies, have enabled this information to be used in newer and more comprehensive ways than originally conceived of and correspondingly, not mentioned in previous work on RCM.

In particular it fuels the shift by the company away from the Static methods of life cycle costing, and towards the Proactive methods of whole-of-life costing. This is a step that enables companies to set up the data capture techniques and practices required to propel it towards the Stochastic, or probabilistic, model of whole of life costing.


Measurement of maintenance performance

Measurement is worth mentioning within this paper because the data that will be generated from applying an effective maintenance program will allow for companies to look further at how their program is functioning than they previously could. It is another example of the fundamental importance of effective maintenance policies.

When applying measurement programs to asset management or maintenance, companies generally look directly to direct performance measures. These are things such as failure rates, mean time to repair, availability, quality and a whole list of other measures of how a machine is operating at any given time. (And often trended to give a view of improvement or deterioration)

These are perfectly reasonable measures and they give a company a snapshot of how a machine is performing to the standards that have been set for it. Regardless of how these measures are selected or generated, they are almost always *lagging* indicators. That is, they are indicators that tell you how your machine is performing after the fact.

Figure 7. Areas covered within the RCM Scorecard



Yet the data collected through establishing an effective maintenance program allows the company to generate a range of **leading** indicators. Measures that lead performance, or tell you that something is likely to begin to perform badly before it actually does.

The diagram in Figure 7 depicts the relative impact of these areas of **leading** indicators, and the smaller impact of performance measures established in the traditional **lagging** approaches. These are the key areas of the RCM Scorecard, a tool first published in the book, *The Maintenance Scorecard*^{xii}, and the subject of a separate article in this area.

However, the basic thrust of the RCM scorecard is to allow companies to measure the effectiveness of their maintenance policy initiatives. Through applying measures to the data captured in the course of doing the day-to-day work, RCM Analysts are able to establish things such as:

- ❑ Is it more cost-effective to manage the asset, over its whole-of-life profile, or not? (**Leading** to incorrect whole-of-life management, not just costs)
- ❑ Was the task really more cost effective than the estimates of failure? (**Leading** to incorrect whole-of-life costs)
- ❑ Was the cost of failure really more cost effective than the estimated costs of the maintenance policy? (**Leading** to incorrect whole-of-life costs)
- ❑ Are the tasks actually predicting or preventing failures?
- ❑ What is the increase in risk due to late performance of D_{Tive} tasks? (**Leading** to higher than acceptable levels of risk exposure)
- ❑ What is the increase in risk due to late performance of P_{Res} or P_{Rep} tasks? (**Leading** to higher than acceptable levels of risk exposure)

- What is the increase in risk due to late performance of P_{Tive} tasks? (*Leading* to higher than acceptable levels of risk exposure)

The actual measures contained within The RCM Scorecard are detailed fully within the book. It provides, arguably, a stronger level of benefit to a company than direct measures because it allows them to tap into the results of mainly leading indicators, thus heading off poor performance before it appears on the management report. Regardless of the actual measures used, the point remains that this is only possible due to the creation, in the first instance, of the effective maintenance program.

The foremost consideration of maintenance managers

We began this paper discussing the three key drivers of maintenance that EAM systems often target. Without considering the different operating environments of different companies, these do cover the basic drivers of most maintenance departments.

1. To develop a maintenance policy designed to minimize the total cost of managing and operating the assets throughout its entire life cycle for a given level of performance and risk.
2. Obtain maximum efficiency out of the resources used to carry out the maintenance policy, driving unit costs further towards the optimum level.
3. To steadily build the asset data portfolio to allow future decisions regarding the asset base with increasing levels of confidence.

It could be argued that methods based in RCM style thinking alone could satisfy all three of these primary drivers of maintenance management. But the business processes that would be required to do so would be onerous and would restrict the ability of the company to manipulate and analyze data effectively. As well as being a burden to those trying to manage the maintenance workload.

It could also be argued that implementing EAM or CMMS, without implementing a parallel, or leading, initiative to create an effective maintenance policies will produce limited results, possibly exacerbate the current situation by allowing the company to perform incorrect work efficiently. And potentially create an environment where the assets are being managed in a way that is contrary to principles of responsible asset stewardship.

This line of thinking can lead to only one conclusion. The development of effective maintenance policies is the **foremost** consideration for modern asset managers. When done correctly it provides the base for business processes, inventory management techniques and methods, software configurations and selection, and the numbers and skill requirements of labor.

Aside from these tactical advantages, it also offers the strategic advantages of improving the whole-of-life management and understanding of the physical asset base, and the way that it is monitored and managed through performance measures. However, once the program is created, attempting to manage the asset base without leveraging off the advances in modern maintenance software deprives the organization of a tremendous opportunity for improvement.

This viewpoint is not new, nor is it particularly complex. It is a common sense approach and is an extension of the basic way that maintenance managers acted prior to discovering technology and being drawn down the path of increasing functionality,

graphing, mobile devices and other gadgetry. It just seems to have been lost in the maze of tools that we are faced with today!

Of particular importance in this paper is the growing role of the RCM Analyst. Once a separated facilitator or a sole analyst, the RCM Analyst is a role that is by necessity becoming a lot broader. Covering a range of additional areas of expertise. A 20th century facilitator was generally driven to apply a team-based method and to complete the analyses. A 21st century Analyst is generally the owner of the program for their area or region. They are responsible for its upkeep, implementation, for ensuring that it is effective, for establishing the links to whole-of-life costing, and for capturing the knowledge of the organization through the application of the method in a flexible fashion. A new role for a new set of challenges!

Daryl Mather is a specialist in the areas of physical asset management. He is the author of the book “The Maintenance Scorecard” and currently works with selected clients on a global basis.

ⁱ MRO stands for Maintain, Repair, and Operate and is an acronym widely used within the EAM/ERP industry and associated with inventory management from an asset perspective rather than from a production perspective. The difference is that with ERP style inventory management the focus is on “just-in-time” methods. While MRO style inventory management focuses on “just-in-case”, or probabilistic methods.

ⁱⁱ The author acknowledges that the definition of what is an acceptable, or unacceptable, standard of performance is an extremely complicated area and one that would take several articles to cover in adequate detail.

ⁱⁱⁱ Within asset management cost-effectiveness is **not** merely low direct costs. Rather the minimum costs for a given level of risk and performance. (Maximum value)

^{iv} The Iowa Division of Labor Services, Occupational health and Safety Bureau, issued a citation and notification of penalty to Cargill Meat Solutions, on the 30th of January of 2006. This citation and notification or penalty required corrective actions such as the establishment of a preventive maintenance program and training of maintenance personnel on potential failure recognition among a range of initiatives to be implemented. This is just one of a number of recent safety events where maintenance has been flagged as a contributing factor.

^v Anecdotal information provided to the author from senior management within a range of companies in the water industry of the United Kingdom places asset failures as being responsible for approximately 40 – 60 % of breaches of consent. In this context “consent,” relates to guidelines designed to protect the environment to an agreed level. In infrastructure this is thought to be even higher. This particular industry represents one of the world’s first water networks and much of the infrastructure is ancient.

^{vi} Mathematical Aspects of Reliability-centered Maintenance, H. L. Resnikov, National Technical Information Service, US Department of commerce, Springfield

^{vii} Mathematical Aspects of Reliability-centered Maintenance

^{viii} GEM stands for generic error modeling and was first developed by Professor Rasmussen of MIT following his review of the incidents leading up to the three-mile island disaster in the USA. The field of human error is a fundamental area of modern reliability management and has been advanced greatly by the works of James Reason, of Manchester Business University in the United Kingdom.

^{ix} Reasonably likely is a term used within the RCM Standard, SAE JA1011, to determine whether failure modes should, or should not, be included within an analysis. “Reasonableness” is defined by the asset owners

^x The strategy options, or policy options, offered within RCM are detailed in the RCM standard SAE JA1011.

^{xi} This could theoretically, be suitable for all companies that need to manage physical assets. However it has particularly importance for financially regulated institutions and companies that need to prove the case for funding.

^{xii} The Maintenance Scorecard, ISBN 0831131810, is published by Industrial Press