PENSPEN

MANAGING THE CONSEQUENCES OF FAILURE

in Asset Integrity

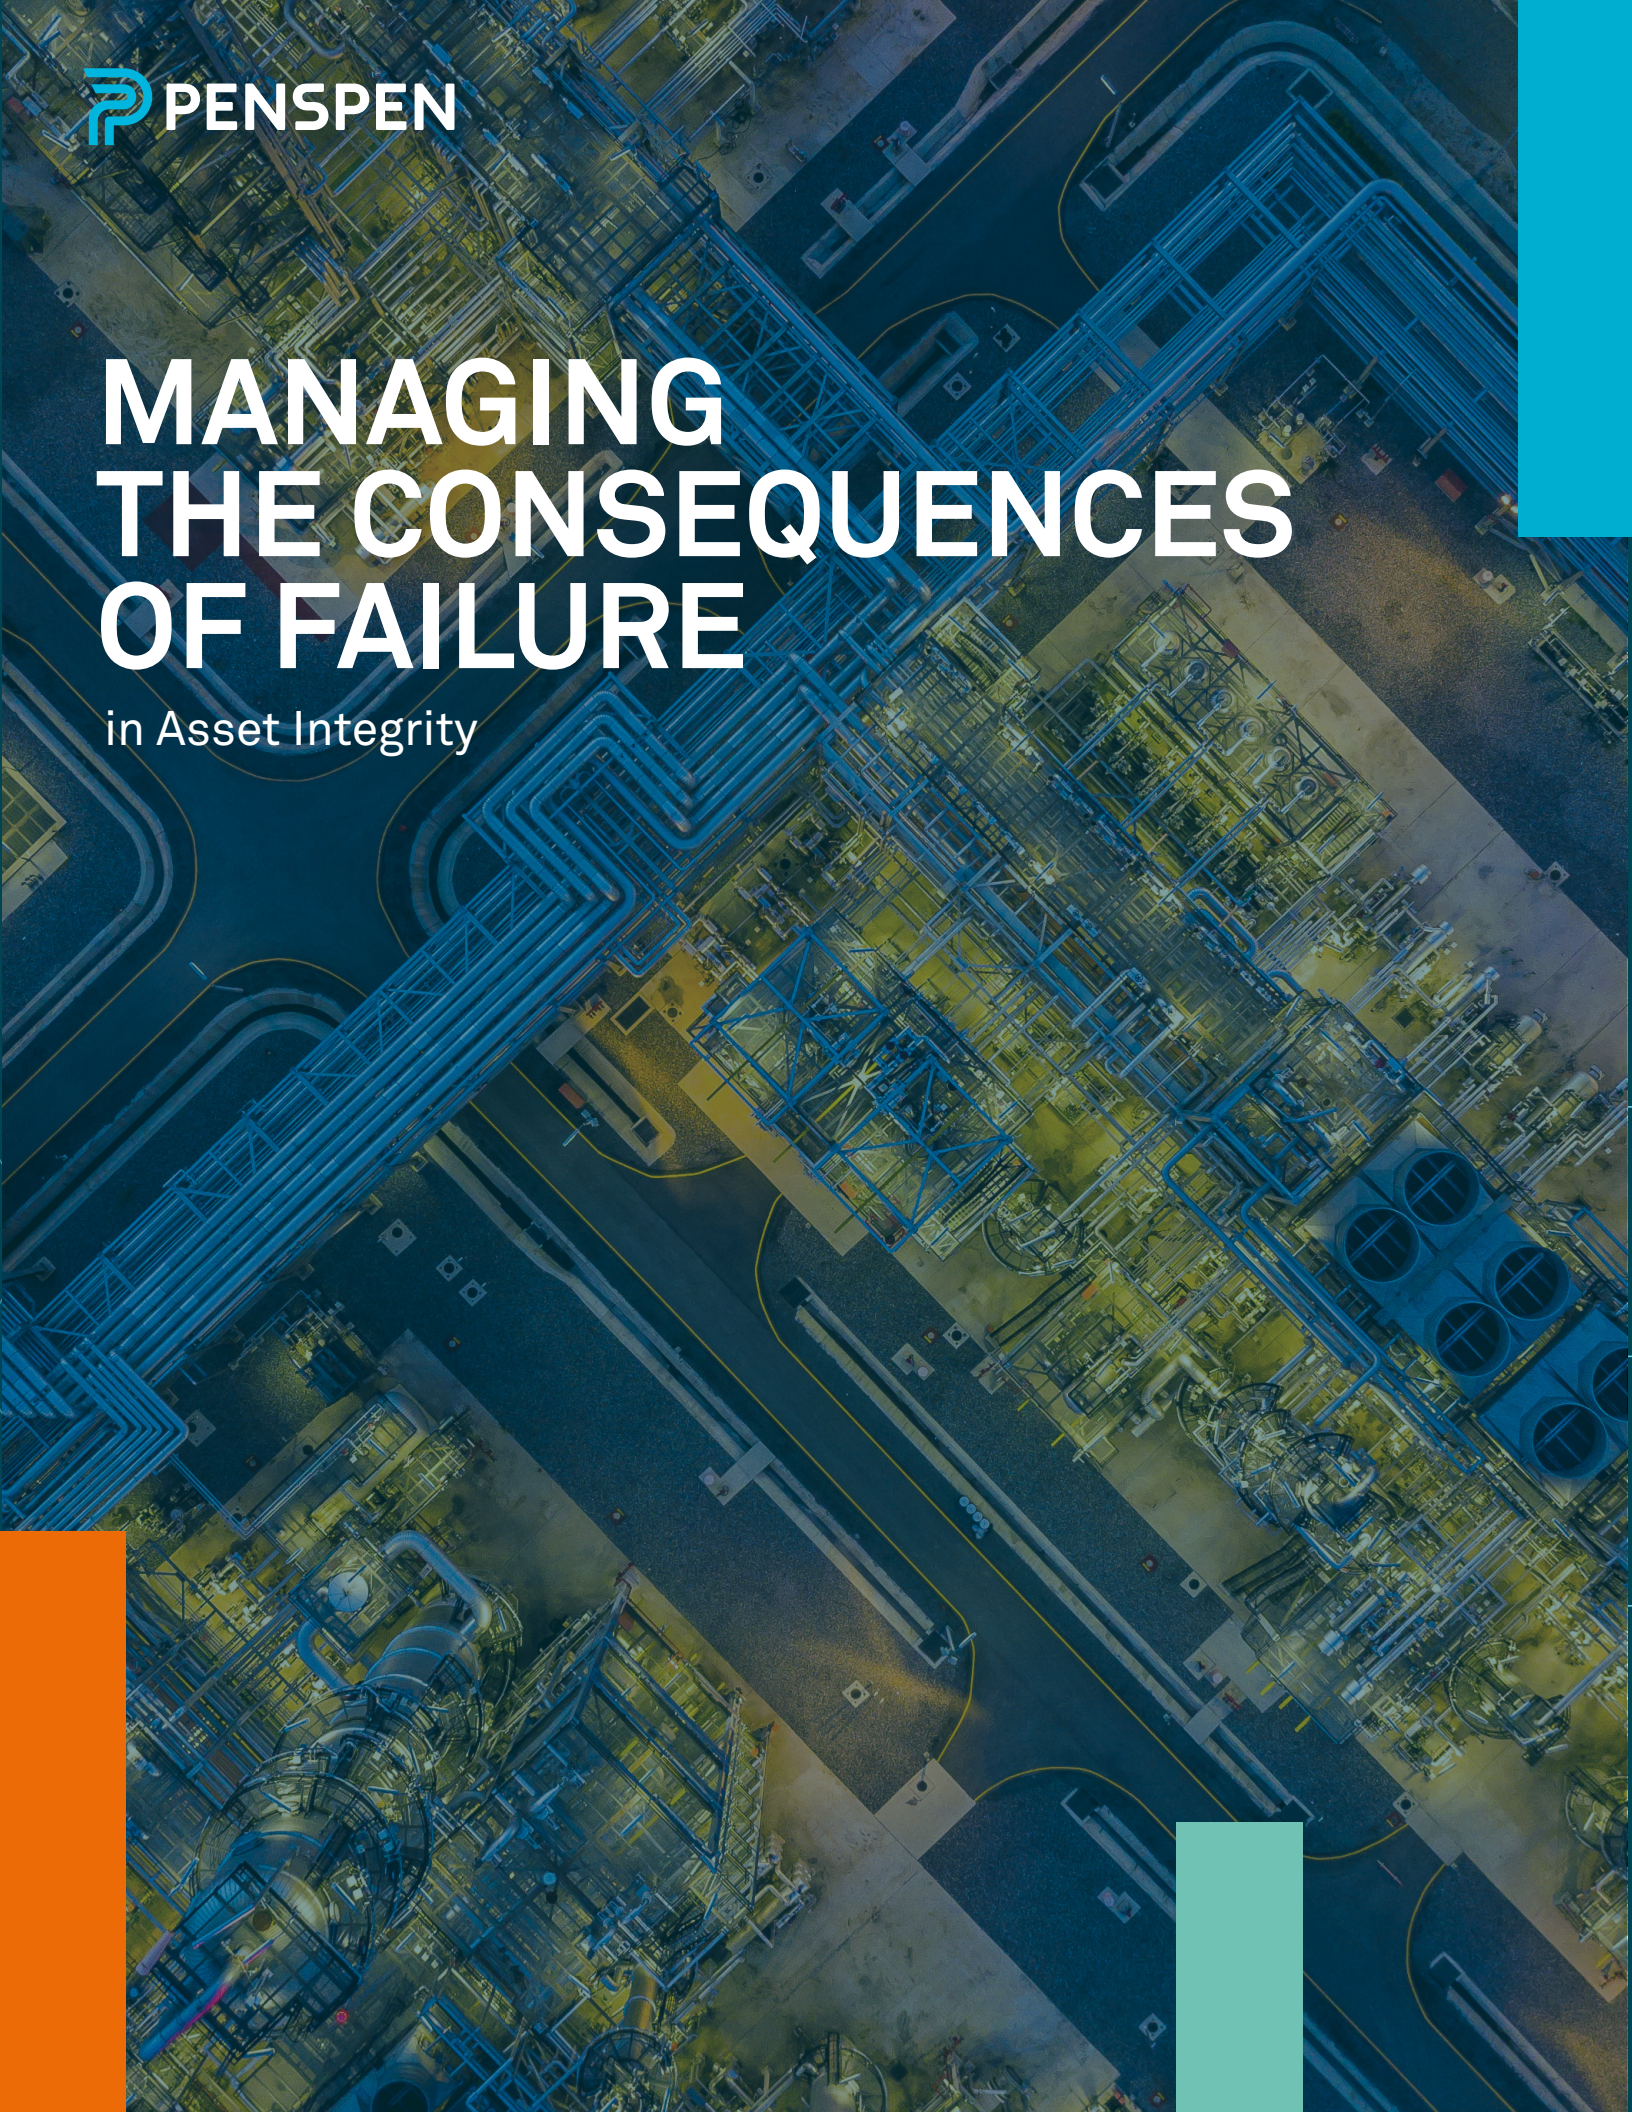

Contents

1. Abstract

With relatively few exceptions, wherever engineering goes maintenance is sure to follow. Inevitably however, so is equipment and/or asset failure together with the accompanying consequences of the failure. For redundant systems and equipment, the ability to maintain the desired level of functionality is achieved by the inclusion of duplicate or multiple components, thereby in the event of a Functional Failure (FF), there is an alternative means of providing the primary or secondary functions.

The same situation however does not pertain to the many standalone assets and systems operated within industries whereby there is no redundancy. Moreover, little or no guidance exists, with end users solely reliant on the maintenance recommendations of the Original Equipment Manufacturer (OEM), modified as necessary to meet statutory and regulatory requirements. However, the OEM is unaware of how the end user is using these assets and systems, many of which have serious consequences in the event of a single Failure Mode (FM) occurring. As such, which OEM would provide three different strategies for identical assets used in three different ways? Could an OEM do so even if they tried?

With maintenance accounting for a significant proportion of overall operating budgets, there is a general inherent ethos that it is regarded as a monetary burden. Instead, maintenance should be seen as ensuring that any physical asset continues to do what the end user wants it do, in its present operating context. Maintenance therefore should be viewed as an investment; a direct cost that organisations should be ready and willing to bear with the expectation that they will receive a benefit ensuring safety compliance; operability; and asset/system reliability that far outweighs the magnitude of any investment.

Maintenance is concerned with the preservation of function, namely assets/systems are procured with the objective of doing something, with maintenance ensuring that they continue to perform to the end users' satisfaction. To derive objective maintenance therefore, it is necessary to know exactly what is required in terms of performance. Accordingly, if the deterioration of any asset or system can be predicted, it may be possible to identify an action to ensure any such deterioration and accompanying consequences can be managed: and, with suitable intervention; be reversed. This is in effect what Reliability Centred Maintenance (RCM) does: identifies the functions, the functional failures, the failure modes, and what can be done to predict such failure modes. Further, RCM also takes account of what happens when things fail and can tailor maintenance requirements depending on the consequences of failure.

How therefore can industries optimize the maintenance strategies of standalone non redundant assets and systems in order to manage the consequences of failure?

This paper aims to provide an insight and guidance into the potential benefits for managing the consequences of failure with respect to non redundant systems through the application of RCM methodology; recognising that the consequences should be of greater concern than the engineering failures themselves, thereby concentrating resources on those failures that actually matter.

2. List of Abbreviations

Abbreviation	Definition
CM	Condition Monitoring
FF	Functional Failure
FM	Failure Mode
FMEA	Failure Modes, and Effects Analysis
FMECA	Failure Modes, Effects and Criticality Analysis
MTBF	Mean Time Between Failures
NSM	No Scheduled Maintenance
OEM	Original Equipment Manufacturer
RCM	Reliability Centred Maintenance

3. Introduction

Expressed as the number of assets required and the total number available; redundancy is the existence of one or more additional or standby assets, which if required; can perform to the desired level of functionality in the event of a failure of the primary asset. The existence of multiple assets will clearly improve the reliability and availability of any system while managing the consequences of the original failure be it safety; operation; or non-operational. However, what if an asset is non-redundant, namely there is no spare? Furthermore, what if, in the event of a failure there is a significant impact on the safety of personnel?

If the consequences are this serious, they require an effective management strategy in order to prevent or reduce the probability of the failure that causes the consequences. In recognising that the consequences of any failure are of greater importance than the engineering failures themselves; resources in the form of maintenance can be targeted where they most matter. This therefore requires a suitable and objective maintenance strategy.

However, in order to derive objective maintenance and tailor requirements accordingly, it is necessary to know exactly what is required in terms of asset performance. If the deterioration of any asset or system can be predicted therefore, it may be possible to identify an action to ensure any such deterioration and accompanying consequences are appropriately managed and; with suitable intervention; reversed.

A proven approach for such a maintenance management strategy is Reliability Centred Maintenance (RCM). It regards maintenance as the means to sustain the functions an end user requires of an asset or system in a defined operating context. Defined as “a process used to determine the maintenance requirements of any physical asset in its operating context”; it has four distinctive features: identification of the functions to be preserved; the functional failures; failure modes; and what can be done to predict such failure modes. Further, RCM also takes account of what happens when things fail tailoring maintenance requirements depending on the consequences of failure.

No single industry is the same, nevertheless assets that may be considered as non redundant include:

- Deck cranes
- Wire lifting rope
- Single engine installations

- Pump/motor sets
- Air compressors

Ideally, the RCM process should be conducted when an asset is first identified during design. However, the main objective of this paper is to provide guidance on the development of a maintenance strategy for in service assets using the RCM process; outlining the various stages of an RCM maintenance deriving study; and the interpretation of seven basic questions. If applied correctly, it can transform the relationship between the organisations that utilise it; their existing assets and systems; and those personnel who operate and maintain those assets and systems. Furthermore, it will also recognise that the consequences of failure should be of greater concern than the failures themselves.

4. Origins of RCM

A report commissioned in 1974 by the US Department of Defence asked United Airlines (UAL) to conduct a survey of maintenance trends in the commercial aviation sector. The results of the survey were produced in a report entitled Reliability-centred Maintenance. Written by two of UAL's senior reliability engineers; Stanley Nowlan, and Howard Heap; the report proposed a decision based maintenance regime founded on managing the consequences of failure rather than attempting to prevent every failure. This view was formed, in part, by the significant amount of redundancy being built into commercial aircraft to ensure their operability.

Moving forward and the work undertaken by Boeing during the 747 development; which had resulted in the Maintenance Steering Group (MSG) 1 and 2 methodologies, the Nowlan and Heap report gave rise to a number of interpretations of RCM in various standards including MSG 3, mainly within the US military. With the production (and demise) of various specifications between the US military arms; by the early 1990's all but the NAVAIR specification remained and was generally regarded as the best in the business, albeit primarily focused on the airline industry.

During the early 1980's however, John Moubray, along with his associates began working with the application of RCM aimed specifically at the mining and manufacturing services, using a slightly modified version to that of Nowlan and Heap. Moving to the UK in the late 1980's, he established a commercial company which licensed the RCM2 methodology to a range of companies providing RCM services, including training.

5. Application of RCM

In order to optimise the benefits of RCM, the process should be applied using a top-down approach. RCM is also a "zero based" process, namely it is undertaken as if nothing is being done to predict, prevent or mitigate for failures that could occur. It assumes that there is no maintenance program in place and no asset spares are available to enable recovery from failure. Furthermore, the analysis process is approached with no pre-conceptions of the required maintenance.

6. Functions and Operating Contexts

To apply RCM consideration must be given to the correct interpretation of seven basic questions about the asset or system being studied:

1. What are the functions and associated performance standards of the asset in its present operating context?

2. In what ways can the asset or system fail to fulfil its functions?

3. What causes each functional failure?

4. What happens when the functional failure occurs?

5. How much does the failure matter?

6. Can anything be done to prevent or predict the failure?

7. What should be done if prediction or prevention is not possible?

The first four questions amount to a functionally based Failure Modes and Effects Analysis (FMEA). In answering question 5, a criticality assessment is undertaken which turns the FMEA into a Failure Modes, Effects and Criticality Analysis (FMECA), and it is these latter three questions that will address the failure consequences, and the proactive maintenance strategy using an appropriate algorithm. The severity however of failure consequences is dependent upon how and where assets and systems are used; namely their operating context.

The Operating Context

The purpose for which a particular asset or system was intended may not be appropriate to the environment or context within which it is being utilized. The operating context is the foundation on which the subsequent RCM decisions are made. Because of this it is imperative that the operating context is produced and fully understood in order to not only successfully apply RCM, but implement any maintenance strategy.

The operating context or functional statement is the circumstances in which a physical asset or system is required to operate. It must be produced first, and distinguishes between what an asset or system 'is', from what the asset or system is there to 'do'; its primary (and secondary) function. After all; if it is not known what is expected of an asset; how will it be known to have failed any expectations?

Such statements represent the objectives of maintenance; which is to ensure that any asset or system continues to perform to the end users' satisfaction. It must therefore be appropriately structured to contain information relevant to the normal mode of operation as an aid to understanding; and to provide a common format for each one produced. The following headings are a guide to the structure of the operating context statements (see annex A for further guidance).

- Functional Output
- System Description
- Modes of Operation
- System Availability
- Analysis Boundaries
- Environmental Conditions
- Redundancy
- Protection
- Hard-Wired Condition Monitoring (CM)
- Pre/Post-use Checks
- On- Receipt Checks
- Despatch Checks
- Failure Mitigation
- Hazardous Material (HAZMAT)
- Assumptions
- Supporting information

Equipment redundancy is a particularly important issue. When it exists, the circumstances when recourse requires to be made to a standby asset or system, or some sort of system reconfiguration must be clear in the operating context. Furthermore, redundancy will generally reduce the overall consequences of failure. However, as the tile of this paper suggests, where there is no redundancy, how can the consequences of failure be managed?

In What Ways can the Asset or System Fail?

In general terms, an asset or system is said to have failed when what it is required to do is outside the boundaries of what it can do, and can thus no longer undertake the defined and measurable performance requirements. Defined as "the inability of any physical asset (or system) to fulfil a function to a standard of performance which is acceptable to the user" (Moubray, J.). It can come about in two ways; a full functional failure or a partial functional failure.

Function	Function Failure
To pump 600 litres of water per hour	• Does not pump any water at all • Does not pump 600 litres of water per hour

The differences between the two different functional failures is that in the first example there is no output at all; whereas in the second functional failure it is still doing something, but below the defined and measurable performance requirements.

The differentiation is necessary, as each functional failure will have different causes, with each of those causes generating different consequences and therefore a different management strategy. As RCM therefore is a process that takes account of the consequences of failure, there needs to be an understanding of the manner in which assets/systems fail. These need sufficient understanding to enable them to be associated with one of the patterns of failure to ensure that assets/systems can achieve their inherent reliability in their specified operating context.

7. The Physics of Failure

More or less since time immemorial the nature of failure has been represented by the traditional (but perhaps now) misconceived bath-tub curve.

The curve is reasonably un-complicated. On the "Y" axis is the conditional probability of failure; the curve plots the number of failures within a population during each period of operation, the expected life occurrences forming the "X" axis. When items are new, many of them succumb early to failure, followed by a period of steady state failure or "random" failure. Eventually the remaining items experience an increased wear out until all have reached the failed state.

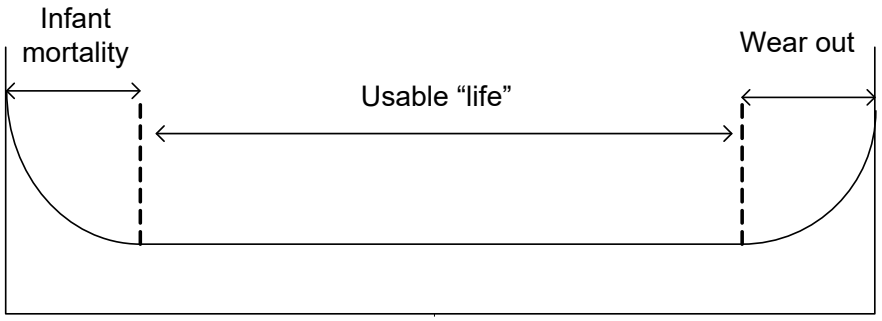

Figure 1: Traditional "bath-tub" Curve

With most engineers likely to have been introduced to this view at some time or another; the question posed therefore, is: if a motor vehicle conformed to this traditional pattern of failure, how long would such an OEM remain in production? Thus if it is not tolerable in personal everyday life, why do many large industry organisations, regulatory and governing societies accept it?

As was stated in the "Origins of RCM" above, the commercial airline industry discovered that maintaining aircraft as though every item had a useful "life" was counterproductive (Nowlan and Heap report). Challenging conventional thinking, it was revealed that the majority of failures in modern complex equipment did not follow the "traditional" representation of failure as had been previously assumed. What was instead established was that there were in fact six failure patterns.

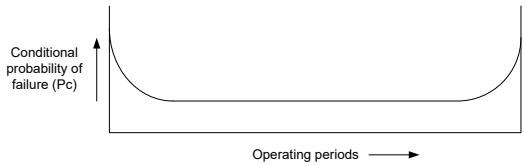

Figure 2: Bath-tub Curve

Pattern A - "bath-tub" (<4%)

Infant mortality, followed by a constant or slowly increasing failure rate, then a distinctive wear out zone.

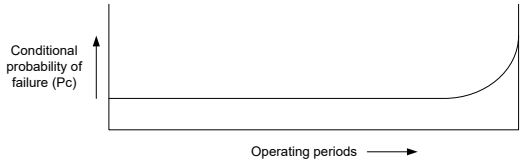

Figure 3: Life Curve

Pattern B – "life" (<2%)

Constant or slowly increasing failure rate followed by a distinctive wear out zone.

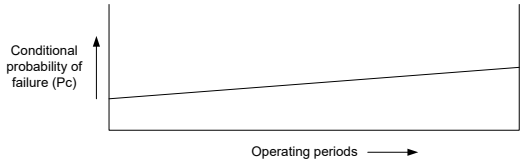

Figure 4: Rising

Pattern C – "rising" (<5%)

Gradually increasing probability of failure, but no distinct wear out zone.

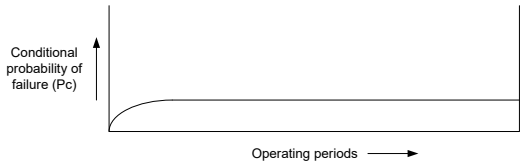

Figure 5: Honeymoon

Pattern D – "honeymoon" (<7%)

Low failure probability initially, then a rapid increase to a constant failure probability.

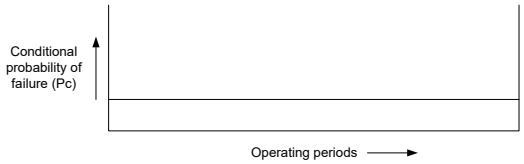

Pattern E – "random" (<14%)

Constant probability of failure in all operating periods.

Figure 6: Random

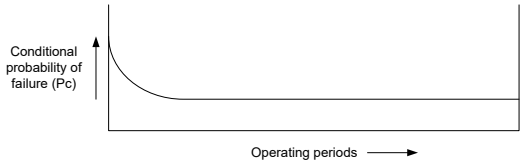

Pattern F – "infant" (>68%)

High infant mortality followed by a constant or slowly rising failure probability.

Figure 7: Infant Mortality

Can any conclusions be drawn from these failure patterns? If the failure patterns are divided into those that generally exhibit "life" or age related characteristics (A, B, C); and those the exhibit "random" characteristics (D, E, F); then the resultant failure statistics become 11% for those items exhibiting "life" characteristics and 89% for items that fail randomly. If the pattern of failure therefore is random, carrying out scheduled overhaul or replacement type maintenance tasks will be ineffective. Instead the deterioration of the asset or system must be detected in service or; wait for it to fail.

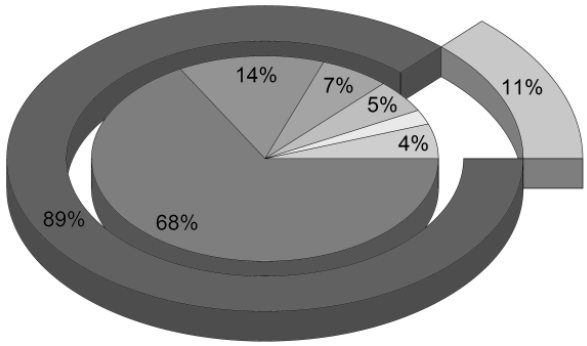

Figure 8: Random and Life Comparison

The key message however from these failure curves is the realisation that if the nature of failure of an item conforms to one curve, but is instead allocated to another, then it is extremely unlikely that an effective failure management strategy will be derived.

An example of this is the perception of a failure mode conforming to failure pattern B in the belief that there is a wear out period following a useful "life"; when instead it conforms to failure pattern F which has a period of "infant mortality". The result will be that following the undertaking of "life-based" maintenance, a period will follow where there is a high conditional probability of failure before the item once again "settles in".

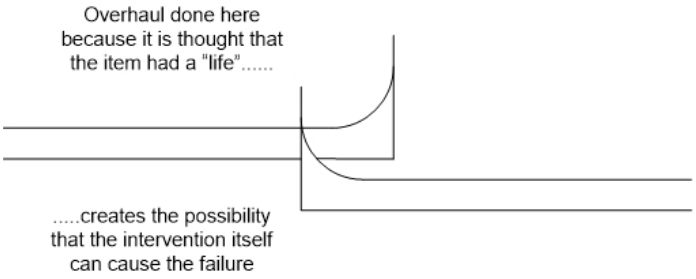

Figure 9: Bath-tub Curve

This phenomenon is frequently observed, and will deliver the worst possible maintenance strategy. The inference therefore is that; unless there is a dominant failure mode, imposing an age limit does little or nothing to improve the overall reliability of a complex item. In fact, in many cases scheduled overhauls actually increase the overall failure rate by "re-setting the clock" and introducing a high infant mortality rate in an otherwise satisfactory asset or system, (as illustrated by pattern F). For this reason, intrusive and intervention type maintenance should be avoided where possible, as what is absolutely right for pattern B... is absolutely wrong for pattern F.

What Causes Each Function Failure?

Failure is distinguished by two criteria. The first is the functional failure; namely the failed state when an asset or system no longer undertakes defined and measurable performance requirements. The second is the event that causes the failed state, the Failure Mode; and can be defined as "a single event that causes a functional failure" (Moubray, J.). They are the link between the functional and physical worlds and as such represent the failure characteristics of the asset or system providing the function(s); they are the physical causes of functional failures. They should be recorded in sufficient detail to enable an appropriate failure management strategy to be identified.

Giving consideration to the pumping of 600 litres of water per hour, what failure modes would cause the pump not to pump at all?

Function	Functional Failure	Failure Mode
To pump 600 litres of water per hour	Does not pump any water at all	• Power Fails • Motor Windings Fail • Motor Bearing Seizes • Pump Seal Fails • Impellor Jams • Inlet Blocks • Outlet Valve Jams Shut • Pump Casing Ruptures • ...etc...

Where do you stop? Listing too few failure modes and/or insufficient detail leads to a superficial analysis and fail to identify those failure modes that result in the more serious consequences. Alternatively, too many failure modes and/or too much detail leads to analysis paralysis. The level of detail at which the cause of failure is identified therefore, depends on many factors, paramount among which is the competency of the vessel/platform operator/maintainer. Furthermore, it is at failure mode level where maintenance is managed (as opposed asset/system level or component level).

In practice however, only those failure modes which might reasonably be expected to occur in the operating context in question are recorded which includes failure modes which have occurred before on the same or similar assets and systems; which are already the subject of proactive maintenance; which have not yet happened, but are considered to be real possibilities. Where the consequences are very severe, then the more unlikely failure modes should also be considered.

Once a comprehensive list of failure modes causing each functional failure has been identified, consideration can then be given to both the effects and consequences in order to implement an effective management strateg

What Happens when the Functional Failures Occur?

There is a distinct relationship between failure modes and their effects. Failure effects are the physical manifestations (if any) that result from the occurrence of a failure mode. They describe what happens when a failure mode occurs where nothing is being done to predict, prevent or mitigate the failure (zero based). However, the narrative must be sufficiently concise to enable the consequences of the failure to be determined, since RCM is predicated on the recognition that the consequences of failure are much more important than the technical characteristics.

Therefore only failures with intolerable consequences need to be prevented. It is recommended that the failure effects be described at three distinct levels and are written chronologically. However, there is no absolute time limit for the whole scenario to come to its conclusion, and could be from days to years and is thus "over a period of time".

Local Effects: Those that occur near the failure mode, i.e. within the compartment or enclosure. It is a "fly on the wall" account of what is seen, heard, smelled or felt.

Higher Next Effects: Those that occur, typically, in a compartment e.g. control room, or an area remote from where the failure is occurring or elsewhere in the system and would include remote indications and alarms.

End Effects: What happens to the platform, installation, vessel, rig etc... and/or to the operating crew and can include:

- Secondary damage
- Project effects
- Repair action
- Repair time
- Spares parts
- Repair costs (inclusive of any secondary damage)

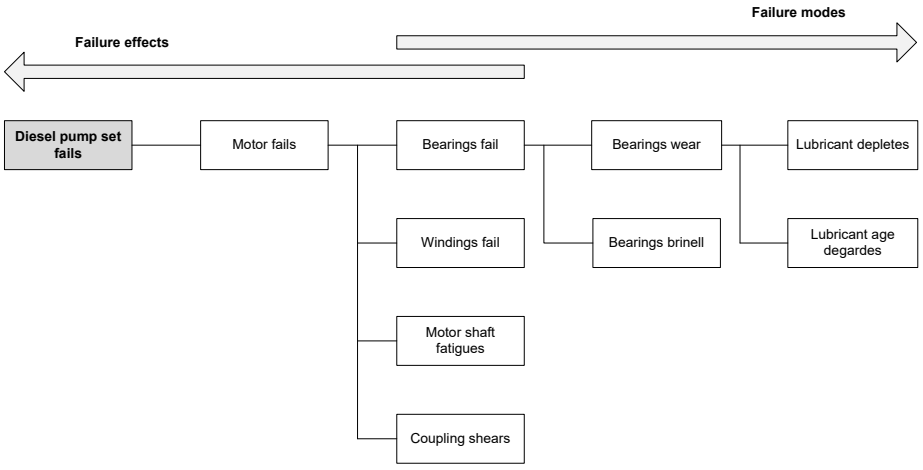

Figure 10: Relationship between failure modes and failure effects

8. The Consquences of Failure: How Much Does it Matter?

In recognition that the consequences of failure should be of greater concern than the engineering failures themselves, how much does any failure matter?

		Severity				
		Negligable	Minor	Moderate	Serious	Major
Probability	Very Likely	Medium	Medium	High	High	High
	Likely	Low	Medium	Medium	High	High
	Possible	Low	Low	Medium	Medium	High
	Unlikely	Low	Low	Low	Medium	Medium
	Very Unlikely	Low	Low	Low	Low	Low

Each time a failure occurs, the user and/or owner of the asset or system is affected in some shape or form, for example; loss of production and increased operating costs. Nonetheless, some failures on their own may at first have no effect but may however eventually create a risk of a more serious failure. As such, if these failures are not prevented, further effects such as resources for repairing which could be better used elsewhere will ensue.

It is the severity and characteristics of these effects that govern the consequences: namely, how much does the failure matter. If the regularity and/or the severity of failures can therefore be reduced, the associated consequences will also be reduced. It should be noted however, that personnel involved in an analysis will need guidance on meaningful equivalents; e.g. how probable is "unlikely"? The issue of probability is often open to interpretation).

With the focus being on the consequences of failure, RCM assesses the effects of each failure mode before classifying it into one of the four general categories of failure consequence: hidden; safety/environmental; operational; and non-operational.

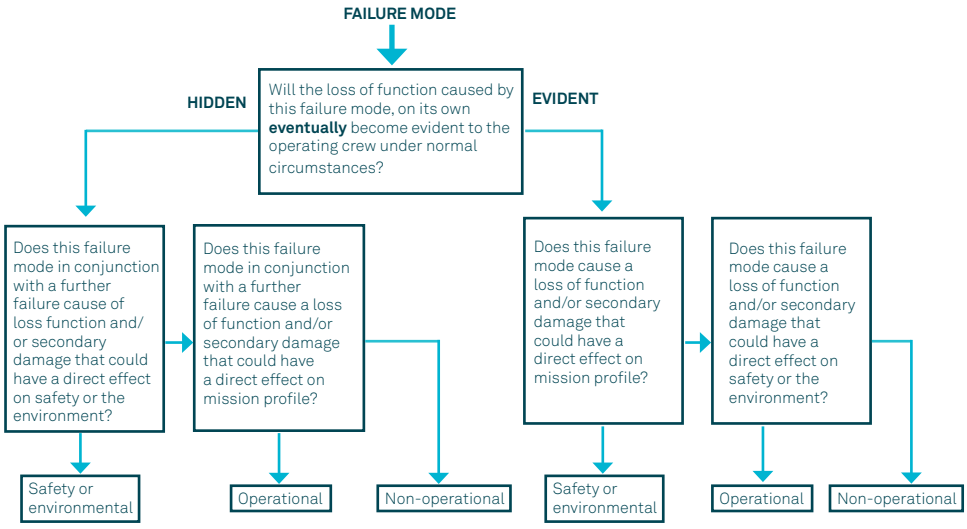

Hidden Failures

Hidden failures generally apply to protective devices which fail in such a way that no-one knows they have failed. Namely it is a failure mode which on its own will not become evident to the operating crew under normal circumstances, whereby nothing is being done to prevent or predict the failure (zero based). "On its own" therefore indicates that there is an issue of a multiple failure.

Safety/Environmental Failures

Safety/environmental failures involve those where there is a significant impact on the safety of personnel; or where an environmental standard may be breached.

Operational Failures

Operational failures have a direct adverse effect on the operational capability of the vessel of platform and can affect operations in number of ways – but by how much? Loss of project or mission profile; reduction in process output i.e. reduced availability and/or downtime; reduction in product quality; reputation; increased operating cost. Furthermore, they all have an economic dimension.

Non-Operation Failures

Non-operation failures only have an effect on the direct cost of repair.
Having classified the consequences of failure, how can they be managed? The next stage in the RCM process is to assess whether it is physically possible to carry out a proactive task that will reduce or, enable an action to be taken that will reduce the consequences of failure to an acceptable level.

9. Managing the Consequences of Failure

On Condition Tasks

The selected task(s) under consideration must be applicable and effective in addressing the failure mode directly, not the effects: as these are the symptoms of failure and so infer that failure has already occurred. As such the technical characteristics of the task must match the technical characteristics of the failure; and reduce the probability of the failure to a tolerable level. Furthermore, it must be worth doing!

RCM is not solely a condition-based maintenance strategy, but gives consideration to all forms of maintenance, and even whether there is a need for maintenance at all. It does nevertheless attempt to move as much preventative maintenance as possible to "on-condition" tasks and requires the identification of suitable condition monitoring (CM) techniques. Such tasks will be suitable where failure modes exhibit measurable deterioration before functional failure which leads to the concept of the progression-to-failure period, or the P – F interval.

The P-F Interval

The P – F interval is used to determine how often an on-condition task should be undertaken whereby the minimum initial task interval is calculated as less than half the P – F interval for safety/environmental consequence failure modes and less than the P – F interval for other consequences.

It should be noted that the P – F interval has almost nothing to do with "life" and absolutely nothing to do with Mean Time Between Failures (MTBF). It is the lead-time to failure – i.e. it is not about when a failure occurs, but how quickly a failure occurs once the deterioration process starts in earnest.

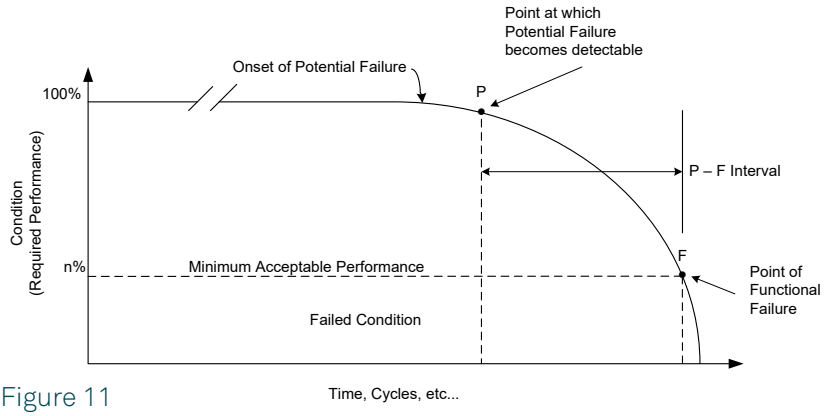

Figure 11

If a potential failure therefore is detectable between point P and point F, it may be possible to carry out an intervening action to either prevent or avoid the consequences of the functional failure. This will be dependent on the detection method used to locate P which will be detectable at different positions on the curve (N.B. the order may differ for different types of failure).

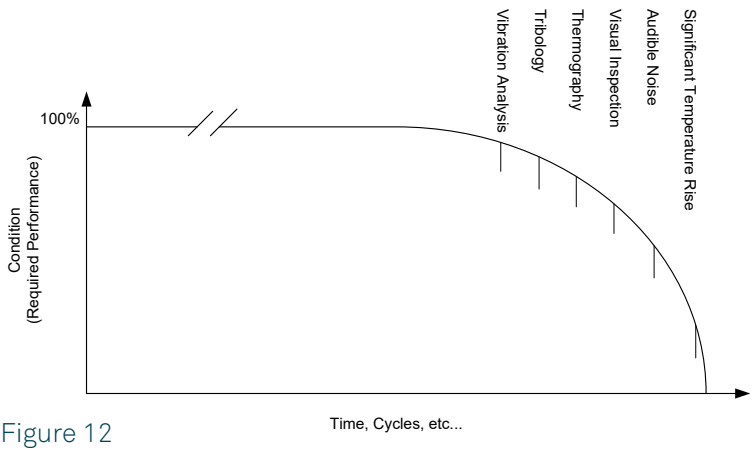

Figure 12

Nevertheless, the potential failure should be a clearly defined and identifiable phenomenon. This means that there must be a detectable, and preferably measurable, change in asset/system condition. The time interval between such phenomena becoming detectable and the functional failure occurring (the P – F interval) must be measurable. However, although the actual length of the P – F Interval for most assets/systems has been observed, it is unlikely to ever have been measured. This information resides in people's heads and has to be extracted and made sense of.

The P – F Interval must also be long enough to be of use for action to be taken to avoid or reduce the consequences of the failure. This means that at the time of detecting the potential failure, there should be sufficient time remaining to plan and set in motion a remedial action before the functional failure occurs. But importantly, the identified on-condition task must be feasible; able to be done at the required interval; worth doing; and be more cost effective over a long period of time than the full cost, including secondary damage, of allowing the failure to occur. However, not all failure modes exhibit measurable deterioration before functional failure.

Hard Time Tasks

Where there is no measurable deterioration before functional failure, on condition tasks will not be applicable or effective. Instead, consideration needs to be given to Scheduled Restoration and Scheduled Discard or "hard-time" tasks (Moubray, J., 1997, pp13 – 14). These type tasks can only be applicable and effective if the failure conforms to patterns A, B, or C and a useable "life" can be determined. For pattern C, where there is no discernible wear-out zone, it is the increase in the conditional probability of failure to a tolerable level that will determine "life".

Within RCM analysis, most "lifed" items can be considered as conforming to pattern B. However, for a hard time task to be selected, it must be shown that the failure mode for the asset/system in question exhibits a distinct and definable life. Further, if a restoration or overhaul task is considered, it must additionally be shown that the asset can be restored to its inherent reliability and/or level of performance. This is the same for discard tasks, where it must also be shown that the required function can be reinstated by replacement of the failed asset/system.

10. What if predicting or preventing the failure mode is not possible?

Where no preventative task can be identified that is both applicable and effective in managing a given failure mode, then a default action, governed by the consequences of the failure should be evaluated. What default tasks represent is a conscious decision to either do nothing (allow the failure to occur) or take some other form of action to modify the consequences of failure, or eliminate the failure mode completely.

No Scheduled Maintenance (NSM)
If no tasks match the criteria for applicability and effectiveness; then No Scheduled Maintenance (NSM) is an option, but only if the consequences are either operational or non-operational. It must not be used for safety or environmental consequences, regardless of whether the loss of function is hidden or evident.

However, opting for NSM does not mean doing nothing. There will still be a requirement to ensure that the associated downtime is minimised, for example by having a spare available. There will also be a requirement for a corrective (remedial) task to be identified.

Failure Finding Tasks
Failure finding tasks are also an option, but only for hidden failures when an applicable and effective preventative task cannot be specified. This is because for the full consequences of a hidden failure to occur, there must be a subsequent failure such as that of a duty item when the protective device (the stand-by) is already in the failed state.

Protective devices can therefore be in the failed state without it being realised until there is an independent but functionally-related failure. When such items are in the failed state, they can only be discovered with a periodic failure finding task. It should also be noted that a failure finding task may find the subject item in the failed state. In this respect, failure finding is a preventative task, in that it prevents the consequence of a multiple failure.

Redesign or Change Action
Changing anything is expensive. However, if a failure has safety or environmental consequences and no applicable and effective task can be identified, then some form of redesign, or one-time change is necessary. The objectives of such actions are the same: to reduce the consequences of failure (perhaps by reducing the probability of failure or eliminating the failure completely). Redesign involves a hardware modification; whereas a one-time-change relates to a modification of operating or maintenance procedures; skills (training); spares holdings; or the imposition of limitations on asset/system performance levels or operating contexts. Whichever course of action is taken, the objectives of redesign and change action must be either to reduce the likelihood of failure or to reduce the severity of their consequences.

11. Summary

The reliability of assets is linked to the maintenance to be applied to them since it will determine not what will need to be done in terms of maintenance (which is defined by the design), but how often a maintenance intervention is likely to be required.

Reliability can be defined as: “The probability that an item will perform a required function without failure under stated conditions for a stated period of time”. Since the conditions must be a constant, given a defined mode of use, unreliable equipment fails before the stated period of time has elapsed.

Quite often this can be a subjective judgement, but reliability is a statistical measure of performance, often expressed in terms of Mean Time Between Failures (MTBF) that can be expressed and evaluated in mathematical formulae dependent upon reliability characteristics.

For the purposes of RCM, the distinction has to be made between two types of reliability characteristics: random failure, where a failure can occur at any time (although there may be a detectable warning period) and wear-out, where failure probability increases with exposure to stress.

RCM is a process that takes account of the consequences of failure. It needs an understanding of the way in which assets fail - their physics of failure. These should be sufficiently understood to enable them to be associated with one of the patterns of failure to ensure that assets can achieve their inherent reliability in their specified operating context. It is used to determine the maintenance requirements of an asset in its specific operating environment to ensure that it continues to achieve its required performance standards. It allocates the most suitable maintenance, with the least expenditure of resources, and recognises that inherent reliability levels cannot be improved upon through maintenance if the original design or mode of operation is inadequate and is not subsequently modified.

Using the wrong maintenance techniques can waste money, time, and resources, while having little or no effect on managing the consequences of any asset failure.

Determination and commitment will be required at all levels of management if RCM is to be introduced against a fixed cultural background to the process, and there must be commitment by management to procure “fit-for-purpose” training for relevant employees.

There must also be participation by all teams for the implementation of an RCM strategy, which must be without exception in order to achieve maximum success; an expectation that must be reinforced to and by management at all levels if the benefits are to be realised. Furthermore, a cultural change is required due to newly ascribed risk of approach toward degrading (not failure of) assets/systems, with an awareness such that RCM is not discredited.

It should also be noted that there is a growing awareness within regulatory and governing societies, including classification societies that sole reliance on OEM recommendations as an approach to maintenance needs to be reviewed. Lloyds Register for example will allow the use of RCM techniques as part of an approved Machinery Planned Maintenance Scheme provided certain criteria are met (Machinery planned maintenance and condition monitoring, March 2013).

However, introducing RCM to a complete platform, installation or site in “one hit” will consume a significant number of resources. It needs to be a structured process which should include a review and evaluation of the most critical assets/systems, namely those whose consequences of failure are the most severe.

RCM is a structured and auditable process which focuses on sustaining outputs. It has been tried and tested for almost four decades by an industry with two key values: a compelling need to ensure that they do the right things for safety, operational and economic reasons; and a commercial desire to eliminate any non-value added maintenance tasks. If implemented correctly, RCM can provide a return on the costs of analyses at any level, be it “savings” or “cost-avoidance”. Furthermore, it contributes to improved safety and environmental protection; provides higher asset/system availability and reliability; resulting in greater maintenance efficiency and cost effectiveness.

12. Bibliography

Anthony, M., Smith P. E., 1993, Reliability-centred Maintenance, ISBN 007059046X ISO 14224

Ministry of Defence, 2012, Defence Standard 00-45 Using Reliability Centred Maintenance to Managing Engineering Failures, Issue 3

Mokashi, A.J., Wang, J., Vermar A.K., 2002, A study of reliability-centred-maintenance in maritime operations

Moubray, J., 1997, Reliability Centred Maintenance, 2nd edition

Nowlan, F. S., Heap, H. F., 1978, Reliability-centred Maintenance, AD-A066579

SAE JA1011 [3], Evaluation Criteria for RCM Processes

13. Appendix A

Functional Output
This is the reason the asset or system was procured, namely what is it required to do? If the output is variable, the worst case scenario should be use for analysis purposes.

System Description
This should be a physical description of the system's assets and their interconnections including the achievable performance of individual assets identified for analysis. The difference between what performance can be achieved and what the actual user requires should be stated. Careful consideration should also be given to specifying any configuration changes.

Modes of Operation
The normal modes of operation of the asset or system should be recorded. This should include statements on:

- How can the system configuration be altered, e.g. cross connections
- Starting and stopping routines including details of auto/local/remote control
- Time spent continuously operational and/or extent of periods of dormancy. For assets normally dormant e.g. fire pump sets, the readiness times should be recorded.
- If there is extended period of shutdown or dormancy for an asset or system, any preparations required should be recorded.

System Availability
The required asset or system availability for the required undertaking that is aligned to a next higher level operating context should be identified.

Analysis Boundaries
The boundaries of each analysis should be defined in terms of:

- The inputs form other assets or systems that have been assumed to be available as required and hence may be analysed separately.
- The systems that are consumers of the outputs from the assets or systems being considered which may also be analysed elsewhere. For both inputs and outputs the physical boundary should be represented where possible on a system diagram which clearly shows the line of segregation.
- The assets and systems within the defined boundary that have been identified for RCM analysis.

Environmental Conditions
The typical conditions under which the asset or system is used, e.g. winter in the North Sea, or alternatively more tropical climates. Where there are significant extremes, alternative analyses may be required.

Redundancy
Are there any standby systems/assets, i.e. those assets that are normally dormant but can be activated to provide the desired function in the event of a functional failure of the primary asset, should be identified.

Protection
Any protective devices present within the system are be identified, together with the device or event they are intended to protect (or protect against), their operating parameter(s) and mode(s) of operation.

Hard-Wired Condition Monitoring (CM)
Any permanently installed monitoring devices that are present, the parameter(s) they measure and what occurs when preset values are exceeded are to be recorded.

Pre/Post-use Checks
When standby or dormant assets are prepared for use, any current pre-use checks that are undertaken shall be recorded (but are not to be automatically templated into an analysis). Similarly, any current post-use checks or operations to prepare assets for a forthcoming period of dormancy/long period of shutdown shall be noted.

On-Receipt Checks
When assets are delivered from stores, warehouses or other sources for use, any on-receipt checks or tests are to be recorded.

Despatch Checks
When assets are disembarked from a vessel or platform, any preparations, checks or tests required prior to being landed are to be recorded.

Failure Mitigation
Any routines that can be invoked by users/operators of the system following partial or total functional failure(s) of the asset(s) being studied, and intended to replace in part or wholly the function produced by the asset experiencing functional failure, other than by the activation of standby assets, should be recorded. (In other words, consequence mitigation by other systems.)

Hazardous Material (HAZMAT)
Any hazardous material present in the asset to be analysed or which could arise from operating or maintaining the asset, e.g. hydrocarbons, POL waste products, etc.

Assumptions
Any assumptions that have been made, e.g. level of manning, current test and trials and operating procedures.

Supporting Information
Glossary of technical terms used in the analysis.
Bibliography – manuals, makers' handbooks, environmental and safety standards, etc., are to be recorded.

14. About the Author

Gordon Philip is Director of Asset Integrity East Region at Penspen. He is responsible for the provision of technical oversight and guidance to Asset Integrity engineering teams, aiding clients to minimise the risk of asset failure.

Gordon has over thirty-five years' engineering experience including consultancy and technical management covering a wide range of disciplines and multiple market/industry sectors. A certified Maintenance and Reliability Professional (RCM), he holds an LLB(Hons), a BEng(Hons), and is a member of the IMechE Safety and Reliability working group.

15. About Penspen

Penspen is a global team of engineers who design, maintain, and optimise energy infrastructure to improve access to energy for communities worldwide. We help meet the world's evolving energy needs by providing consulting, project, and engineering solutions across the entire energy asset lifecycle.

For over 70 years, our teams have delivered more than 15,000 projects to in excess of 100 countries. By helping countries access lower carbon fuels and by extending the useful life of existing energy infrastructure, we help to bring cleaner energy to millions of people in thousands of communities across the Middle East, Africa, Asia, Europe, the UK, and the US.

Penspen is a proud member of SIDARA a leading, privately-owned professional services group with award-winning impact and global reach. As an engineering, architectural, and planning consultancy that values specialty expertise, Sidara is united by a commitment to providing clients with multidisciplinary solutions rooted in quality, innovation, collaboration, sustainability, and technology to deliver social and community impact.

CONTACT OUR ASSET INTEGRITY TEAM

Email: contact@penspen.com | Website: www.penspen.com