RCM 151 · COURSE HANDOUT

Reliability Centred Maintenance (RCM)

A practitioner's guide to functions, failure modes, consequences and maintenance task selection.

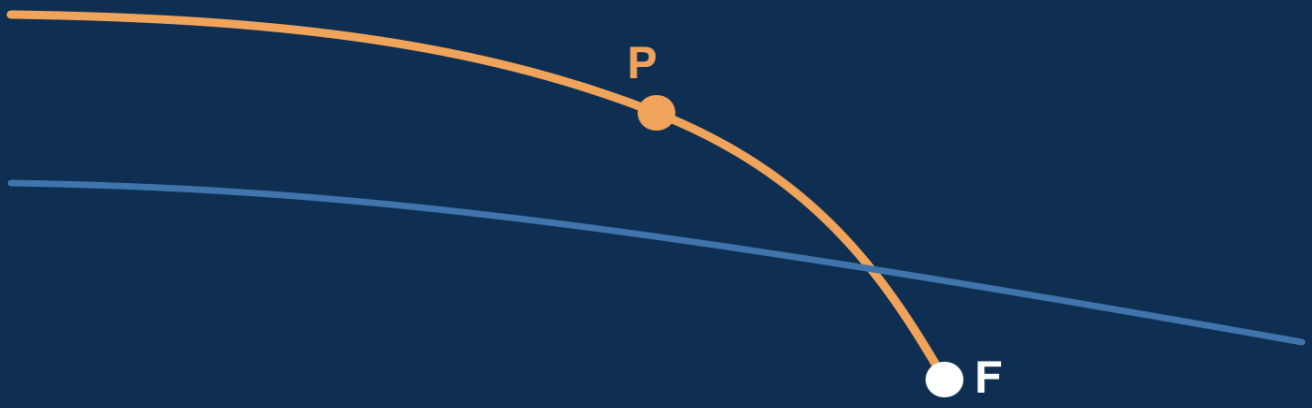

Chris James

Contents

CHAPTER 1

Introduction

1.1 Maintenance Evolution

Prior to World War II, industries weren't highly mechanized and they did not rely much on their physical assets, so equipment downtime wasn't a major concern. Therefore, managers didn't prioritize the prevention of equipment failure. In addition, most machinery was simple, over-designed, reliable, and easy to repair. Consequently, there was no need for systematic maintenance beyond basic cleaning, servicing, and lubrication. The demand for skilled workers was also lower, and people were readily available to handle carry out repairs. Customers, in turn, were more patient.

Preventive Maintenance (PM) emerged just before World War II. Wartime pressures increased the demand for goods, including weapons, while reducing the supply of industrial manpower. This led to greater mechanization and this trend continued after the war. During the 1950s, machines became more complex and numerous, and industries began to rely on them more heavily. As this dependence grew, equipment downtime became a major concern. This gave rise to the idea that equipment failures could and should be prevented. This in turn led to the concept of preventive maintenance.

In the 1960s, preventive maintenance - equipment overhauls and replacements at fixed intervals - was widely used in industry. Regular plant shutdowns took place, with replacement of items such as filters, hoses, bearings, pumps, valves, bushings, seals etc etc.

However, maintenance costs also began to rise sharply relative to other operating costs and pressure rose for the development of alternative maintenance techniques. Unlike traditional preventive maintenance strategies, which rely on intervention at fixed frequencies, Condition Based Maintenance (CBM) focuses on monitoring the actual condition of the equipment to determine when intervention is required. This means that unnecessary interventions – when the equipment is still in a good condition – can be avoided.

The concept of condition-based maintenance can be traced back to the 1960s and 1970s when industries began to recognize the limitations of time-based maintenance strategies. Initial CBM techniques to detect equipment degradation included visual inspections and basic vibration analysis.

The development of CBM was significantly accelerated by advancements in technology. This included a wide range of condition monitoring techniques, which use specialized equipment to monitor the condition of other equipment. In the late 20th century, the introduction of microprocessors and digital signal processing allowed for the collection and analysis of large amounts of data from various equipment sensors. This enabled more automated monitoring of equipment health to be carried out.

A further CBM development was the emergence of the Internet of Things (IoT). By connecting industrial equipment to a network of sensors and devices, this allows for real-time data collection and remote monitoring. More recently. advances in machine learning and artificial intelligence have also played a crucial role in refining CBM techniques by enabling the

development of predictive algorithms that can identify patterns and trends in equipment data.

The evolution and development of maintenance techniques described so far means there are three distinct types of maintenance:

1. Corrective (also known as Breakdown Maintenance, Reactive Maintenance, or Run-to-failure): This type of maintenance is performed after a failure has occurred, such as fixing a broken conveyor belt in a manufacturing plant or repairing a malfunctioning air conditioning unit in an office building.
2. Preventive (involving overhauls, replacements, or restoration): This is a proactive approach which involves overhaul or replacement of parts at a fixed time basis to prevent failures from occurring. For example, regular oil changes and tyre rotations on a fleet of delivery trucks or replacing worn-out parts on an assembly line before they cause a breakdown.
3. Condition based (Predictive, On-condition): These are also proactive approaches and include a wide range of techniques to check equipment condition and look for warnings that the equipment is failing or about to fail. This allows interventions to be made before the failure occurs in order to avoid the consequences of the failure Condition-based tasks include condition checks by people, the use of specialist condition monitoring equipment and the use of built-in sensors and machine learning algorithms. For instance, a wind turbine's sensors could detect early signs of wear on its components, allowing maintenance crews to replace them before a major failure occurs, or a heat exchanger may be cleaned because its efficiency (determined by built in sensors and algorithms) has declined significantly.

Depending on the circumstances and the context, the older techniques still have their place, and modern maintenance is a mixture of various techniques. One of the biggest challenges facing maintenance professionals is not only learning about these new techniques but also figuring out which ones are best for their organizations, and in which circumstances. Making the right choices can lead to improved asset performance and cost reductions. However, making the wrong choices can create new problems or make existing problems worse and lead to the worst of all worlds – high maintenance costs and high equipment downtime.

As businesses have become more reliant on technology and automated systems, expectations for maintenance have also grown.

Downtime has always been a concern, but with Just-in-time processes and integrated supply chains, mean that even equipment failures with short downtimes can have a significant impact on operations. For instance, an e-commerce company relying on automated warehouse systems could face severe delays in order fulfillment if a critical piece of equipment breaks down unexpectedly at the wrong time.

The rise of automation has made high reliability and availability crucial requirements in a wide range of industries. Greater automation also means that more failures can affect the quality of products and services. For example, equipment failures in a hospital's diagnostic imaging department could lead to delayed or inaccurate test results, impacting patient care.

Moreover, many failures now have serious safety or environmental consequences, making it essential for organizations to comply with safety and environmental expectations. For instance, a chemical plant must ensure that equipment failure does not cause leaks or accidents that could harm employees or the surrounding environment. In extreme cases, this dependency on physical assets therefore goes beyond cost and becomes a matter of organizational survival.

As our reliance on physical assets grows, so too do their operating and ownership costs. To maximize return on investment, assets must be kept working efficiently for as long as needed. Additionally, maintenance costs continue to rise, making it a top priority for cost control.

These heightened expectations are further complicated by reductions in manpower and overtime, reorganizations, outsourcing, and reliance on the knowledge and skills of an aging workforce.

> All industries face similar challenges. However, in terms of its effective management of physical assets, one industry is generally recognized as being the "best of the best". This industry has developed and deployed techniques that can be applied equally to all industrial sectors.

The industry in question and the development of one of these techniques is described in the next section.

1.2 The development of Reliability Centred Maintenance (RCM)

> The Civil Aviation Industry is widely acknowledged for its outstanding management of physical assets, including exceptional reliability performance of modern aircraft.
>
> However, their performance has not always been exceptional, and they did not get there by chance.
>
> During the late 1950s, the Civil Aviation Industry planned massive expansion, with the goal of carrying more passengers over longer distances using larger aircraft like the Boeing 747 (or Jumbo Jet). One major problem was that the industry's reliability record at the time posed a significant threat to this growth. In the 1950s, aircraft were not as reliable as they are today, and the Fatal Accident Rate (FAR) (a measure of the frequency of aircraft crashes resulting in fatalities) was 60 fatal crashes per million takeoffs. These failures were primarily caused by equipment failure, and this rate was felt to be too high. This led to the consensus that reliability needed to improve, and a big part of this improvement could be achieved by better maintenance.
>
> One key aspect of the way the aviation industry does business is traceability or auditability. For instance, if an engine blade fails, there is a record of who fitted the blade, who made it, and even which mine the metal came from. This audit trail allows for learning from mistakes and applying that knowledge in the future, a crucial element of continuous improvement. To achieve continuous improvement, you need an audit trail to refer back to – you need to have something written down. The feeling in the 1950s was that a robust audit trail should apply to aircraft maintenance programmes. They felt they needed a scientific approach to maintenance – an auditable approach to determine which maintenance should be carried out on aircraft.

In 1958 the Federal Aviation Act was passed and this stated that:

“All Commercial Carriers will institute time limitations, or standards for determining time limitations, for overhauls, inspections and checks of airframes, engines, propellers, appliances and emergency equipment”

This meant that companies needed to not only specify the tasks and their frequency but also the reasoning and logic behind them – they needed a scientific, logic, auditable and defensible approach to design their maintenance programmes.

When examining the reasons for conducting maintenance at the time, they were often found to be less objective and scientific than desired, including reasons like:

- "Because the manual says so…"
- "Because it's always been done this way…"
- "Because it's about time…"

And once the maintenance intervention had started, "While I was there…"

During the 1960s, aircraft industry maintenance was mainly preventive. They took their aircraft apart at fixed intervals and rebuilt them. Their thinking was therefore that the primary reason for maintenance interventions was the belief that there was a "Right Age" when maintenance was necessary to ensure future safety and reliability. In other words, the idea was that components were more likely to fail at a certain point, as illustrated in the graph below:

Figure 1: The Life curve

Their preventive maintenance programmes were about stepping in and replacing and overhauling components before they wore out and became less reliable. They felt that all components behaved in this way, so this approach was applied to virtually the whole aircraft. One problem with a new aircraft is of course to determine when to step in because there are no or few data about failures. In the aircraft industry the collection of data resulting from failures is even more of a problem since failures (which include crashes) are necessary to collect the data which users believe is needed to decide how the failures are to be prevented. This is known as "The Resnikoff Conundrum." This states that we need to have failures in order to have failure data. But critical failures are not tolerable since they may cause death and injury. Therefore, the maintenance program for a critical item must be designed without the benefit of actual failure data for that item.

Clearly it is also necessary to have a maintenance programme in place from the outset – it is not acceptable to start flying an aircraft without a plan in place to maintain it. So the designers of aircraft maintenance programmes in the 1960s had to get the best information available and use experience, judgment and available data about similar equipment to estimate component lives and determine when to step in and overhaul and replace components to prevent failures.

An implication in the above model is that all maintenance is good, and the key to improved reliability is more frequent and more thorough maintenance. When failures did occur, the natural reaction was to do more maintenance – because of the uncertainty, perhaps the increase in unreliability had occurred earlier, and the solution was more frequent and more thorough interventions.

Over time the FAA became increasingly concerned because evidence was suggesting that reliability could not be controlled using either of the conventional ways of modifying the timing of the maintenance or the amount of maintenance. Perhaps the model was not true? Or perhaps there is a flaw in the "Right Age" model?

This caused the CEO of United Airlines (W.C.Mentzer) to ask in 1960:

"Do we understand the fundamental principles which underly the way we maintain our airplanes?" and "Do we know why we do what we do?"

This led to some of the most important reliability research ever undertaken. The Joint Industry Reliability Program (JIRP) cost $100 Million in the 1960s and 70s involved the co-operation of Operators, Maintainers and Manufacturers, under Regulator direction. Its objective was to explore the relationship between equipment age and reliability. When these relationships were researched, there were surprising discoveries relating to the effect of scheduled maintenance on complex items and the downside of scheduled maintenance - this research provided exhaustive proof that in many cases reliability wasn't related to "age" at all!

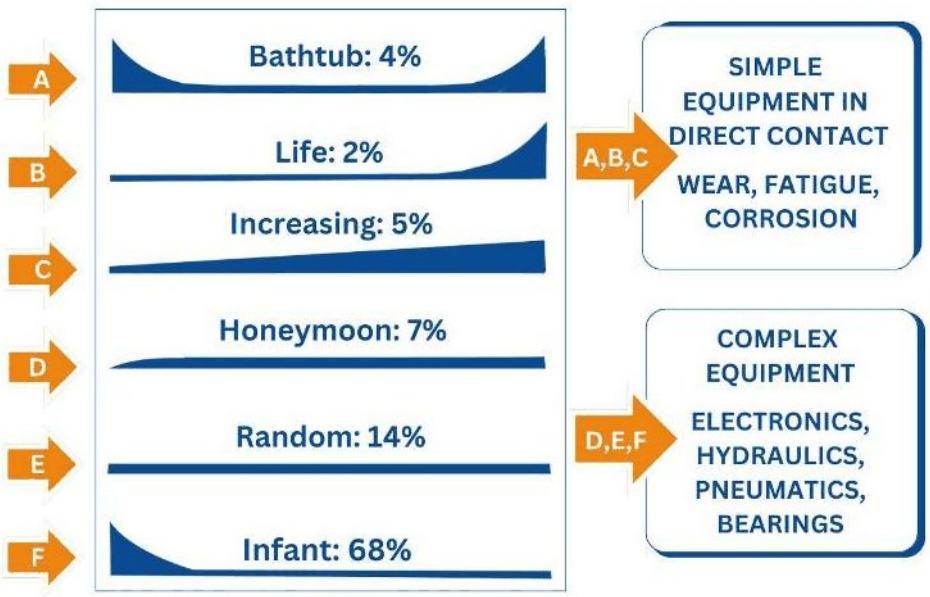

Figure 2: The First Insight – 6 Patterns of Failure

This research revealed that in reality there are six failure patterns. It also showed that 68% of components are more likely to fail because the equipment has been interfered with, or because it is new. These are probability curves that show the likelihood of failure against time for different components.

Pattern A is the well-known Bathtub curve. Research of a large number of components in the aircraft industry showed that this exists in 4 % of all cases. An initial high failure rate settles into a constant failure rate then finally an increasing failure rate. The bathtub curve is

effectively the addition of several root causes. If these are treated separately, they will have different failure patterns. This is why the bath-tub curve is relatively rare at the component level.

Failure Pattern B, where equipment is relatively reliable and then it reaches a point where the risk of failure starts to increase Life patterns was found in 2 % of cases. Mechanisms such as wear, corrosion and fatigue may exhibit pattern B.

Failure Pattern C, with a steadily increasing probability, existed in 5 % of components. This could also be a manifestation of fatigue.

Failure Pattern D, the 'Honeymoon' pattern in which there was an initial increase followed by random failure, existed in 7 % of cases.

Failure Pattern E, Random pattern existed in 14 % of all components. These are components where the probability of failure is independent of the age of the item - it fails 'when it feels like it'. This could apply where failures could occur due to a multitude of unpredictable external factors where failures could occur due to a range of issues in complex systems, and failures where human error is a contributing factor. Another example is most rolling element bearings.

The final pattern is Failure Pattern F - Infant mortality which existed in 68% of all cases. This can be seen in electronic devices. A high initial failure rate might be due to manufacturing or installation defects. Pattern F might also apply to failures caused by human error - in operations and in maintenance.

The fact that different components behave in different ways shows that we need to break failures down into individual causes or failure modes and manage based on their characteristics. Also the high proportion of pattern F shows that many failures are caused by interference, so in most cases, the best tasks are those that minimise interference in a system that is working well. However, it is important to be aware that if these patterns are generated at component or Failure Mode level (see section 3.2.1) and that the existence of Pattern A in this sense may be rare.

The real concern in the industry at the time was that their maintenance programmes assumed everything followed pattern B. But Patterns B and F are mirror images.

Figure 3: Pattern B and Pattern F are mirror images

What's absolutely right for B is absolutely wrong for F! Preventive maintenance is traditionally built around pattern B. But typically only 1 in 50 failures were pattern B, whereas 2 out of 3 failures are pattern F.

Further discussion of Pattern F

Examples of pattern F include human errors due to maintenance or operation, manufacturing defects, design defects and failures due to transport and procurement issues. Data analysis often demonstrates pattern F. For instance, the graphs below show failure data for

piston engine aircraft engines and demonstrate that these engines suffer the highest risk of catastrophic failure when they're fresh out of the factory or a field-overhaul shop. These histograms are derived from data on 180 engine-failure accidents for the five-year period from 2001 through to 2005.

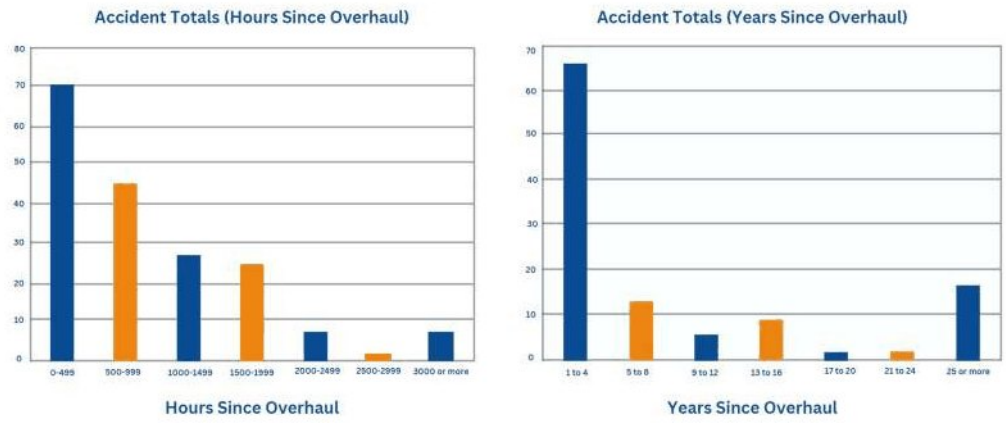

Figure 4: Small piston airplane accidents in 2001 through 2005 attributed by the NTSB to engine

failure, by hours (top) and years (bottom) since engine overhaul* * Source: http://www.avweb.com/news/savvyaviator/savvy_aviator_48_reliability-centered_maintenance_part_2_195969-1.html The implications of pattern F are that whenever there is interference - including preventive maintenance and changes to design - there is risk. Based on the evidence of pattern F, one might conclude that it would be better to do no maintenance, but obviously this is not a viable option either!

Further discussion of Pattern E

Pattern E is a notable failure pattern and represents Random Failure. Components that fail at random fail 'when they feel like it' – there is no relationship between how old they are and how likely they are to fail. In other words, when the component is fitted, you don't know how long it is going to last.

Why do some components fail at random?

Some reasons for random failure are covered below:

- Manufacturing defects: Imperfections or flaws in materials, components, or assembly processes can lead to random failures. These defects might not be immediately apparent and can cause unexpected equipment breakdowns during operation. (Although some manufacturing defects can result in pattern F)
- Material properties: Variability in material properties, such as strength, ductility, and fatigue resistance, can lead to random failures. Microscopic flaws or inclusions in materials can act as stress concentrators, causing premature failure under certain conditions.
- Environmental factors: External influences, such as temperature fluctuations, humidity, corrosion, or vibration, can cause random failures in equipment. These factors can lead to material degradation, component wear, or electrical issues, resulting in unexpected breakdowns.
- Design limitations: Inadequate or suboptimal design can contribute to random failures in equipment. This might include insufficient safety margins, poorly designed interfaces, or

a lack of redundancy in critical components, which can result in sudden equipment failure under certain conditions.

- Human error: Mistakes made during the installation, operation, or maintenance of equipment can lead to random failures (where they are not manifested in early life). This might include improper assembly, incorrect settings, or inadequate maintenance practices.
- Software bugs: Software systems integrated within equipment can experience random failures due to coding errors or unforeseen interactions between different software modules. These failures can result in system crashes, data corruption, or unexpected equipment behaviour.
- Wear and tear: While wear and tear generally cause equipment to fail in a predictable manner, sometimes random failures can occur due to sudden or accelerated wear under specific conditions or due to undetected issues in the equipment.
- Unforeseen interactions: Equipment may fail at random due to unexpected interactions between components, materials, or operating conditions. This can be especially challenging to predict and manage in complex systems with multiple interconnected components.

Which components/ assets might exhibit random failure characteristics?

Further to the above, some examples of random failures are given below:

- Electronic components: Electronic components, such as capacitors, resistors, and transistors, can experience random failures due to, internal stresses, or external influences like voltage spikes or temperature fluctuations. (Note that they can also fail due to manufacturing defects, but this would be exhibited as Pattern F)
- Software bugs: See above.
- Bearings: In theory, bearing failure, due to mechanisms such as fatigue, is predictable, so bearings exhibit age related characteristics. However, in practice, most rolling element bearings fail at random because of variations in real world experience. These variations include load, lubrication, fitting, treatment before fitting, variations in material and variations in manufacturing. (In fact, bearings are one of the few industrial components where a lot of data has been collected and these show that, in most cases, they do indeed fail at random.)
- Any complex component (or a simple component with complex variations in its history). The same logic for bearings could be applied to many other items that do not have dominant wearout characteristics. This includes pneumatics, hydraulics, and a wide range of electrical equipment such as transformers, motors and switchgear. Also, mechanical equipment such as valves and complex mechanical assemblies.

Can anything be done to manage a component that fails at Random?

If a component fails at random, replacing or overhauling it at a given point – traditional preventive maintenance - adds no value. This is because the new components will be no more reliable in the next 1/2/5/10 years than the components than the components that were previously in place. (Since the intervention would have cost money and taken time, overall such an intervention is detrimental.)

This might lead some to think that there is nothing that can be done to manage random failures. However, there often is. Possible activities might include:

1. Activities to reduce the likelihood of failure (by improving the design, the environment or the number of external factors)
2. Activities to reduce the consequences of the failure (in other words, to change the way the failure matters). An example would be some sort of contingency, redundancy or spares
3. Predictive maintenance

The best example of the 3rd Point is a bearing. Most rolling element bearings fail at random – when the bearing is fitted, we don't know how long it is going to last. Therefore, preventive maintenance - replacing or overhauling it at a given point – is no use. However, even though we don't know when the bearing will fail, when it does fail it gives warnings beforehand. Some of these warnings precede the failure by a decent length of time. The principle of predictive or condition-based maintenance is to look for these warnings (which are known as potential failures in RCM) and take action to avoid the consequences of the failure if these warnings are identified.

For a bearing there may be numerous potential failures including noise, heat, vibration, debris in the oil etc.

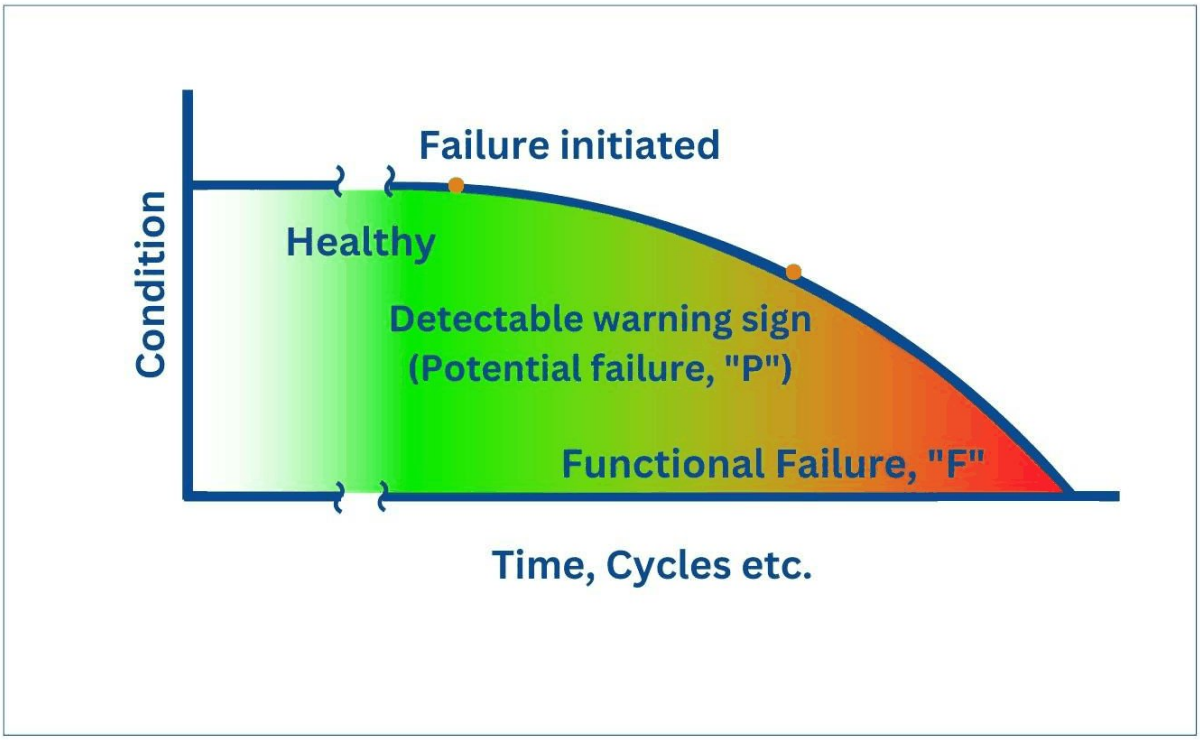

Condition based maintenance is covered in more detail later in this book.

What changed as a result of the research into failure patterns?

The Civil Aviation Industry designed a comprehensive framework for developing maintenance strategies and a result of the 1958 Federal Aviation Act. This framework is known within aviation as Maintenance Steering Group 3 (MSG3), and outside it as Reliability-centered Maintenance, or RCM.

RCM is both a philosophy and a process. The process itself consists of a series of steps. However, there are good reasons for each of the steps (ie thinking and research behind the process). Much of this thinking was informed by the findings from the research into the relationship between equipment age and the likelihood of failure.

Some of these principles that are built into the RCM process include:

- To start from a zero base (i.e. assume no maintenance) when setting maintenance strategies, because interference can sometimes be risky.
- To favour maintenance tasks that are less invasive when selecting maintenance tasks
- To favour on-condition maintenance (which is almost always less invasive than preventive maintenance and can be applied to more failure modes (because it can be applied whatever the failure pattern))
- To consider the management of the consequences of failure as a failure management strategy because 100% of failures cannot be prevented
- To break down equipment failures into individual causes of failure (ie failure modes) and manage each on its own individual merits, because different components fail in different ways, so need to be managed differently
- To recognise that planned maintenance will achieve nothing (in fact it might even be counter-productive) unless correctly targeted
- To recognise that "While you're there..." should be expunged from the maintenance mind-set.

The results of the application of MSG3/ RCM in the civil aviation industry have been dramatic and include:

- Reduced overhauls – on the DC8 (traditionally derived maintenance) 339 items were subject to fixed interval overhaul while on the DC10 (RCM derived maintenance) only 7 items were subject to fixed interval overhaul
- Reduced effort – on the DC8 (traditionally derived maintenance) and 747 (RCM derived maintenance) both have major interventions at 20,000 flying hours, but the resources required are dramatically different o DC8 – 4 million man-hours o 747 – only 66,000 man-hours
- Massive improvements in safety o The overall fatality rate for the industry in 1958 was 60 fatal crashes per million take-offs while in 1998 it was 1 fatal crash per million take-offs o Equipment Failure as Root Cause in 1958 was 40 per million take-offs (2/3 of the total) while in 1998 it was 1 in 30 million take-offs (1/30 of the total or 1/1200 of the historical rate)

Further examples of the benefits of RCM are given in Chapter 8.

1.3 The Four Insights

This section describes some further 'insights', or principles behind the RCM process. An overview of the process is presented in the next section.

1.3.1 The First Insight – 6 patterns of failure

The 6 failures patterns are the first one of the insights in modern maintenance because they provide a clearer understanding of a complex situation and provide a philosophical underpinning of the RCM process.

1.3.2 The Second Insight – the consequence insight

Because maintenance resources are stretched, failures which matter a great deal deserve to have lots of energy expended trying to predict or prevent them, while failures with trivial consequences do not merit the same treatment – in other words, just because you could prevent something from failing doesn't mean that you should. This is also called the Consequence Insight and is best described by considering 3 physically identical assets in a different operating context. Because the consequences of failure are different it is quite conceivable that for certain Functional Failures, the maintenance programme could well be different, even though the assets are physically identical.

We will consider 3 fans of similar age, specification, materials, manufacture etc.

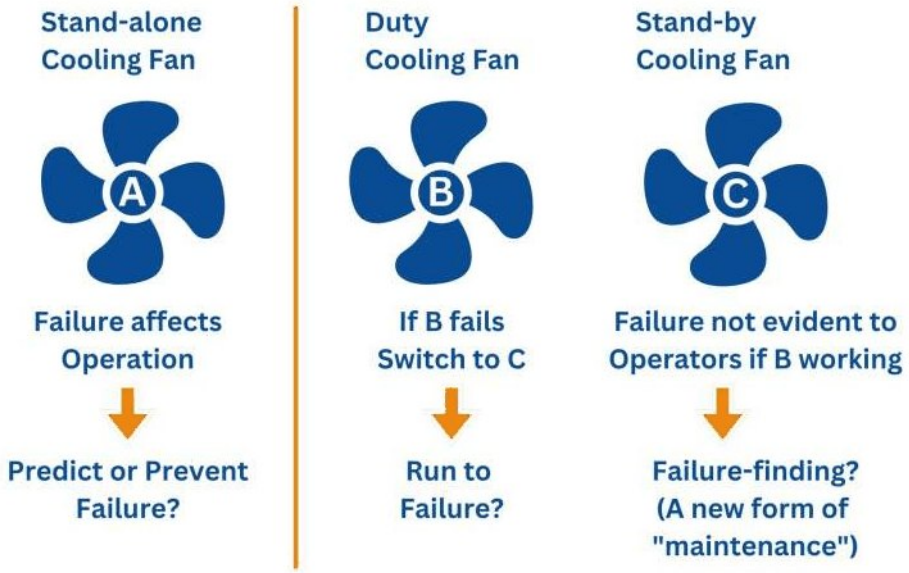

Figure 5: Standalone versus Duty/Standby

A corollary of the consequence insight is the notion of a fourth type of maintenance - Failure Finding Maintenance. This will be addressed later in this document and is a major element of the maintenance of protective devices.

There are significant implications of the second insight. Most preventive maintenance systems are based on the original manufacturer's manual. However, manufacturers do not give 3 different strategies for the same kit used 3 different ways and could not do so even if they wished to because they do not know the operating context of the equipment. This means that we need to look at the operating context of our equipment, in order to decide how to maintain it.

The airline industry realized that they needed a way of looking at exactly how they operated and maintained their equipment. Since such a process did not exist, they had to invent one and went right back to first principles to redesign a maintenance programme.

1.3.3 The Third Insight – what the asset can do and what we want it to do

When the civil aviation industry considered the reasons for maintenance, they realized that maintenance is only necessary to manage failure and that failure only occurs when an asset cannot do what its users want it to do. This led to the realization that any asset has 2 levels of performance – what it can do, and what its users want it to do.

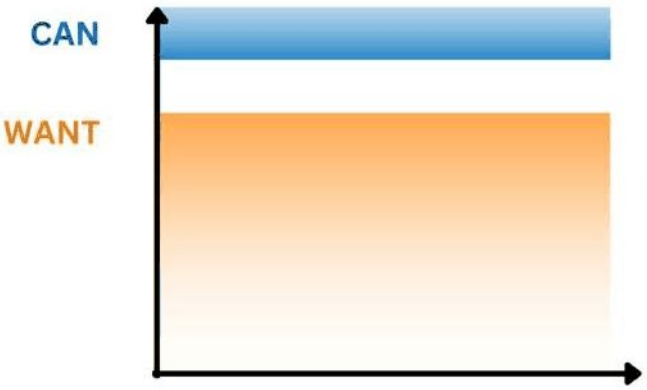

Figure 6: Can versus Want

An asset is failed when what its users want it to do is greater than what it can do. However, reliability problems can be caused by many forces:

- The asset not being "fit" to do the task from the start
- 'Can' falling below 'want' due to deterioration
- 'Want' rising above 'can' during operation

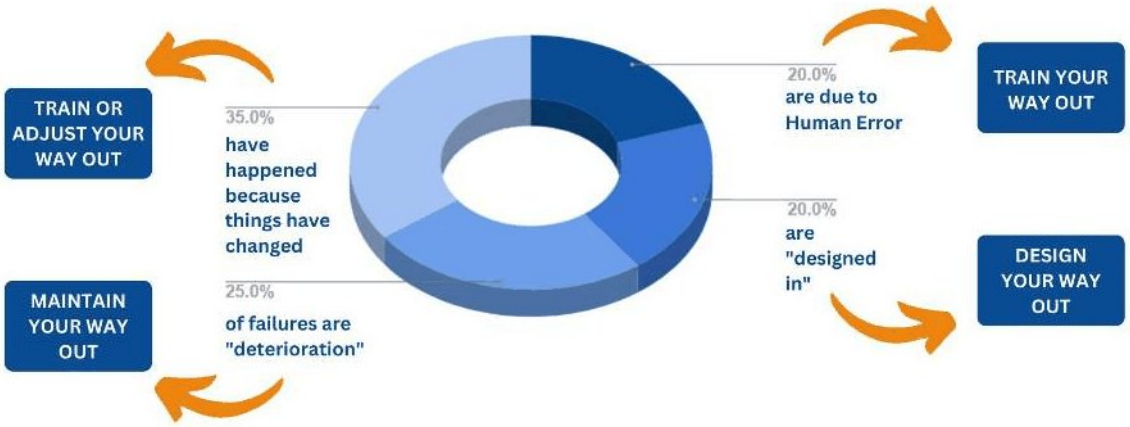

- Some form of human error causes 'want' to be greater than 'can'

Empirical Observation (over the past decade) suggests these forces are all significant contributors to failure.

Figure 7: Four forces contributing to failure

This means that reliability problems require a mix of therapies, but first you must reveal the underlying causes.

1.3.4 Fourth – Team Review

It is clear from the third insight that many people in an organisation are responsible for failures and this must mean that many people in an organisation are responsible for pre-

venting failures. However, tradition has it that "Engineering" can come up with quick solutions to reliability problems. Therefore, we need to harness everybody's intelligence / experience, do so in a concerted, focused, structured way, document it and keep it alive.

This means that the RCM process must be applied by a team of people who know the asset best, as shown below. The benefits of such an approach include Validity, Ownership, Empowerment, Team building and a greater efficiency of review – all of which are requirements for a successful and enduring maintenance programme. (Note that job titles may vary from industry to industry, although the titles shown in the diagram below are typical to many process industries.)

Figure 8: Typical RCM team composition

1.4 RCM Process overview

RCM is a methodology which determines what must be done to ensure that any asset we own or operate continues to do what the users want it to do, in its present operating context.

The RCM Process consists of 7 master questions:

1. What do we want the asset to do? – Functions
2. How can it fail? – Functional Failures
3. What causes the Functional Failures? – Failure Modes
4. What happens when a failure occurs? – Failure Effects
5. How much does each failure matter? – Consequences
6. Can we predict or prevent failure and should we be doing so? – Proactive Tasks
7. How should we manage the failure if prediction or prevention is not an option? – Default Tasks

The first 4 questions form a Failure Modes and Effects Analysis (or FMEA). The remaining 3 questions are addressed by applying a Decision Diagram. This is applied to each Failure Mode and allows the most appropriate failure management decision to be selected for each Failure Mode.

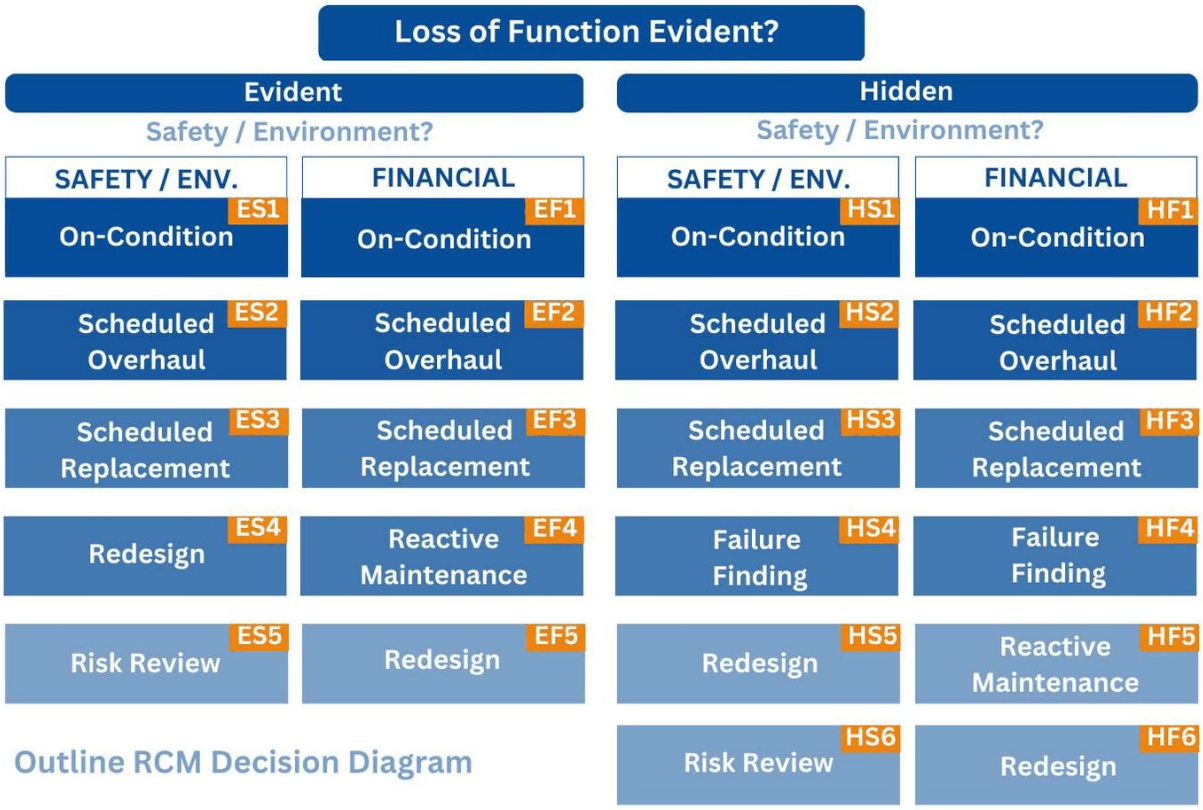

Figure 9: Outline RCM Decision Diagram

1.5 The early results of RCM in the aircraft industry

RCM is applied extensively throughout the civil aviation industry and has been done so for several decades. Back in 1978, one of the architects of the RCM process, F Stanley Nowlan, stated that:

"Cost reductions of this magnitude are of obvious importance to businesses which operate (eets of complex equipment furthermore, such improvements have been made without any compromise on safety."

"A better understanding of the failure processes suffered by modern equipment has allowed us to improve our reliability by focusing our preventive maintenance on speci)c evidence of potential failure"

RCM has been tried and tested for over 50 years by an industry with two key values:

- A compelling need to ensure that they do the right things for safety, operational and economic reasons
- A cold-blooded commercial desire to eliminate any non-value-adding tasks

It has worked in that environment. If your industry has similar pressures and challenges, it can work in your environment too.

CHAPTER 2

Operating Context and Functions

RCM is a process to determine what needs to be done to ensure a physical asset continues to do what its users want it to do in its present operating context. The 7 master questions of RCM are given below.

1. What do we want the asset to do? – Functions
2. How can it fail? – Functional Failures
3. What causes the Functional Failures? – Failure Modes
4. What happens when a failure occurs? – Failure Effects
5. How much does each failure matter? – Consequences
6. Can we predict or prevent failure and should we be doing so? – Proactive Tasks
7. How should we manage the failure if prediction or prevention is not an option? – Default Tasks

This chapter covers the Operating Context and first of the seven questions: Functions.

2.1 The Operating Context

In RCM the operating context is an important concept as it relates to the circumstances and conditions under which a system or component operates. The Operating Context contains any information that informs how the asset fails and the failure decisions to manage each failure mode that is reasonably likely.

It includes factors such as:

- Plant Level Information; what the plant of facility produces and what activities are carried out
- Equipment design and operation: what does the asset under review actually consist of. Manufacturer, model numbers and design characteristics (at a reasonably high level). When it was installed. The quality of the initial installation (if applicable). Any changes since installation. Does it operate automatically or are human interventions needed to operate.
- The Process; The upstream system that feeds the asset and the downstream system that the asset feeds. that may be affected by the failure of the asset being analysed and the level of process instrumentation. This also includes whether there are standby assets, the time before equipment downtime becomes lost income.
- Operational Demands and Windows: This includes the load placed on the system, the frequency of use, the duty cycle (how long it's in use versus at rest), the severity of service, the nature, consistency, and cleanliness of the process materials. In addition, the extent to which demands on the system have changed over time (such as higher output, tighter quality standards, change of supplier, higher quality standards, tighter safety requirements or new environmental legislation). In some environments there are defined shutdown windows, or periods of low demand, that may apply to all assets.

- Existing Maintenance: In general terms! (the objective of the study is to review the existing maintenance programmes). This could include where the maintenance programme came from.
- Environmental conditions: This includes temperature, humidity, exposure to water, dust, vibration etc.
- Regulatory requirements: This includes any standards or laws that the system must comply with, for example, safety or environmental regulations.
- Other factors: For instance, the quality and training of the maintenance personnel (including specialised Condition Monitoring skills), the availability of spare parts, the availability of people.
- Analysis context; Boundaries of the analysis – what is included and what is not. The objectives of the analysis and the current measure of performance in relation to those objectives. Assumptions relating to labour cost and cost of downtime.

Why does the Operating Context Matter?

The operating context matters because it directly impacts the Failure Modes that a system might experience, the effects and consequences of those failures and the management of those failures. In other words, every following stage of the analysis! Understanding the operating context helps to determine the most effective and efficient tasks to manage failures. It also helps in planning and scheduling maintenance tasks, in designing for reliability and maintainability, and in the management of assets throughout their life cycles. Without a clear understanding of the operating context, maintenance efforts may be misdirected or insufficient, leading to unnecessary failures, downtime, and costs.

For example, a pump that is used intermittently in a clean, indoor environment may have different Failure Modes and require different maintenance than the same type of pump used continuously in a dirty, outdoor environment. Similarly, a fan that is protected by a standby may have different maintenance requirements than a standalone fan in a similar process.

Another example of the importance of context is a gas turbine. The same core equipment used to power an aircraft can also be used to drive generators for power transmission onshore or offshore, drive other large items of rotating equipment, or power ships. The context (e.g. load profile, air pressure, air quality, temperature, temperature changes, spare capacity, availability of people and spares, accessibility of the equipment, consequences of failure, time until consequences are felt, etc) will vary dramatically between these different applications. As a result, the failure modes and failure management requirements will be different.

Figure 10: Gas Turbine technology deployed in different operating contexts

How is the Operating Context recorded?

Relevant information is recorded by writing an Operating Context statement. This is a brief (typically 2-3 page) document that includes a general description of the asset(s) or system (at a reasonably high level), and the information described above.

The context statement may also include a high-level system diagram, the boundaries of the analysis and the objectives of the analysis.

Two examples of the Operating Context – a Distribution Transformer and a Central Heating system – are given below.

2.1.1 Example of an Operating Context; Distribution Transformer

The Distribution Transformer is an essential component of the town's electrical distribution system. It is an oil-filled, three-phase transformer, which steps down the high-voltage electricity received from the transmission lines to a lower voltage suitable for distribution to homes, businesses, and other electrical loads in the area.

Typical models are ABB Model: TRM630, TRM710, TRM800. The General Specifications are given below:

- Rated Power: 630 kVA, 710 kVA, 800 kVA
- Primary Voltage: 20 kV or 30 kV
- Secondary Voltage: 400 V or 690 V
- Frequency: 50 Hz
- Insulation Level: 36 kV (HV), 0.433 kV (LV)
- Cooling: ONAN (oil natural, air natural)
- Tank Type: Corrugated fin
- Standards: EN/IEC 60076-1, 60076-2, 60076-3

Dimensions and Weight:

- TRM630: 1770mm x 910mm x 1290mm, 1000 kg
- TRM710: 1770mm x 910mm x 1290mm, 1100 kg
- TRM800: 1870mm x 1020mm x 1385mm, 1200 kg

Other Features:

- Tap changer: Off-circuit
- Winding material: Copper

In the worst case demand the failure of the Distribution Transformer can result in power outages, disrupting the normal functioning of homes and businesses. Some failures can also cause damage to electrical equipment and appliances, and in some cases, create safety hazards.

The Distribution Transformer has been in operation for several years, and its customer base has grown over time. As the town has expanded, new buildings, homes, and businesses have been connected to the electrical distribution system, increasing the load on the transformer.

The Distribution Transformer operates in a typical European urban environment, which includes a mix of residential, commercial, and industrial areas. The climate in the region is moderate, with mild winters and warm summers. However, the transformer may be exposed, on occasions, to extreme weather events such as heavy rain, snow, or high winds.

The operation of the Distribution Transformer is subject to strict safety requirements to ensure the protection of both the electrical system and the public. This includes the installation of proper grounding and protective equipment, and compliance with relevant safety regulations and standards.

The maintenance program for the distribution transformer includes regular inspections and testing to ensure its proper functioning. This includes visual inspections of the equipment, as well as testing of electrical connections, insulation resistance, and oil quality.

The analysis excludes the tap changer which is analyzed separately.

2.1.2 Example of an Operating Context; Domestic Central Heating System

A central heating system is an essential component of many family homes, providing warmth and comfort throughout the colder months.

The central heating system is operated in a temperate climate with occasional extreme temperatures, but normally ranging from 0°C to 25°C. The system is used primarily in the winter months to provide warmth and comfort to the home's occupants. The central heating system is installed in a medium-sized family home, consisting of several bedrooms, bathrooms, living spaces, and a kitchen. The home is occupied by several family members who have varying comfort preferences.

The system is used to provide heat to the entire home. The temperature can be set in each room to a comfortable level based on the home's occupants' preferences. The system is used for several hours each day, with the heating requirements varying depending on the time of day and occupancy. The central heating system is fueled by natural gas, which is supplied to the home via a utility line. The gas is used to power a boiler, which heats water which is pumped to radiators throughout the house.

The boiler is a Worcester Bosch Greenstar 8000 30kW condensing combination boiler. This also heats the hot water to the house on demand.

If the heating system fails, the occupants of the house will be uncomfortable if the outside temperature is cold. However, it is assumed that none of the occupants are vulnerable enough to face a safety risk in the event of failure of the heating system.

The boiler is maintained according to the manufacturer's recommendations, which includes regular inspections, cleaning, and filter replacements. The system is also checked for leaks and other issues to ensure its safe and efficient operation. The central heating system is equipped with various safety features, including temperature and pressure sensors, carbon monoxide detectors, and emergency shut-off switches. The system is also installed according to local building codes and safety regulations to ensure its safe operation. These are also covered by the analysis. The hot water supply is not covered by the analysis.

2.2 Functions

What is a Function?

A Function is a statement of what the users of the asset want it to do

Why do we de)ne functions in RCM?

We only own and operate assets because we want them to do something. If they continue to do what we want them to, for as long as we want, we should be satisfied. Therefore, the development of a failure management programme should start with the consideration of what the users of an asset want the asset to do – in other words, its functions.

Starting with functions is quite fundamental because it goes back to the definition of failure. Dictionaries and standards, may have subtle variations but the definitions of a failure are very similar, and are usually something like:

- The inability to of a system or item of equipment to do what its users want it to do/ to meet an expectation/ to fulfil a function/ to perform something required

Since defining Functions is fundamental to the management of failures. It is logical for a process about failure management to start with Functions.

Which people in the organization are usually in the best position to de)ne functions?

> The people who use the resources, usually those in Operations, often know best how each resource benefits the organization's physical and financial health. So, it's crucial they participate in the Reliability Centered Maintenance (RCM) process from the start.

Do assets have more than one function?

Virtually every asset has many functions. Assets are acquired for a particular reason(s), but there are other requirements that come along as well. The main reason(s) the assets were acquired are known as the Primary Function(s). Additional requirements are known as Secondary Functions. These categories are covered in more detail below.

2.2.1 Primary Function(s)

Primary Function(s) describe why the asset was obtained in the first place. They are the main reason the asset exists. This category often covers requirements such as speed, output, carrying or storage capacity, product quality and customer service.

For example, The Primary Function of a car might be stated as:

- "To transport up to four people from A to B, at speeds of up to 120 kph"

The Primary Function of a pump might be stated as:

- "To pump a minimum of 300 litres/ minute of water into the tank"

2.2.2 Secondary Functions

Secondary functions are other requirements that are not the main reason an asset exists.

> Users also have additional requirements in other ways. These can include safety, following environmental rules and regulations, structural strength, comfort, containment, control, protection, operational efficiency and an acceptable level or appearance.

For example, as a car user, we require other functionality in addition to the primary function written above:

- "To be able to stop the car from 40 kph to rest within 40m, on a dry road in a straight line"
- "To prevent the occupant of the front seat from coming into contact with any equipment or fascia in the event of a collision"
- "To look acceptable"

Additional (secondary) functions for the pump might include:

- To contain process fluid
- To indicate delivery flow to an accuracy of +/- 1%

These functions are not necessarily less important than the Primary Function, they are simply not part of the Primary Function and are additional requirements that need to be maintained (even though they are not the main reason the asset was purchased in the first place).

Secondary Functions may relate to various categories. In fact, a car, has secondary functions relating to all categories of secondary functions. Examples of each are given below:

CATEGORY	EXAMPLE FUNCTION	EQUIPMENT
SAFETY	To prevent the occupant of the front seat from coming into contact with any equipment or fascia in the event of a collision.	Seatbelt and Airbags
ENVIRONMENT	To limit discharge of Carbon Dioxide in the exhaust to less than 170g per km	Engine and Engine management system
COMFORT	To maintain the internal temperature of the passenger cell between 15ºC and 25°C	Climate control system
STRUCTURAL	To be able to withstand a side impact of 25mph without injurious deformation of the passenger cell	Door pillars, floor pan, side impact protection system
CONTAINMENT	To contain the fuel within the system	Fuel tank, filler hose and cap, fuel lines, injection system, manifold etc
CONTROL	To regulate and maintain vehicle speed between 0 – 90 mph	Accelerator pedal and cable, Engine management system, Cruise control?
INDICATION	To indicate clearly the driver's intention to execute a manoeuvre	Indicators, Column stalk, Cabling, Sender unit
PROTECTION	To alert the driver if the oil pressure falls to 2.5 barg	Oil Pressure Warning System (including pressure sensor, cabling, microprocessor, warning lights / readouts / diagnostics)

Figure 11: Table of Example Functions

2.2.3 Protective Devices

Protective devices are an important subset of secondary functions. Some of the different ways they work, with examples, are given below:

1. Preventing Equipment Damage: Some protective devices are designed to prevent damage to equipment in the event of abnormal operating conditions. For example, circuit breakers and fuses are designed to interrupt electrical current in the event of an overload or short circuit, preventing potential damage to the electrical system.
2. Safety of Personnel: Some protective devices also play a vital role in ensuring the safety of personnel. For instance, pressure relief valves in a boiler system protect against overpressure conditions that could lead to an explosion, potentially harming people in the vicinity.
3. Maintaining Operational Continuity by changing the consequences of failure: Some protective devices function to maintain operational continuity by preventing unnecessary shutdowns. For example, redundant systems or backup devices can take over if a primary device fails, allowing operations to continue with minimal disruption
4. Warning of abnormal conditions; This would include alarms and trips on parameters such as high or low pressure, temperature, flow, differential pressure, position, level etc. Also smoke, flame and fire detectors
5. Prevent dangerous situations arising: This would include warning signs

Many protective devices fail in a hidden or unrevealed manner. They may sit in a failed state, and operating personnel may be unaware that they have failed, until they are needed. Then it is too late, and a (relatively) small problem becomes a much bigger one.

Modern plants may have a large number of such devices, and the correct management of these devices is a big challenge for safe and reliable plant operation. The first steps in the correct management of these devices is to identify them, and define how they work and what they need to do. In other words, to define their functions.

Functions Mnemonics (A Mnemonic is a device such as a pattern of letters, ideas, or associations that assists in remembering something.) Mnemonics often used as a reminder for categories of secondary functions. One example is 'SPICES' where:

S; Structural, Safety P; Protection I; Indication E; Environment, efficiency, economy S; Superfluous

This mnemonic does not include the letter A (Appearance). Another mnemonic sometimes used is ESCAPES (although this lacks the letter I (indication)!

Are primary functions more important than secondary functions?

It is not always the case that the loss of the primary function is more important that the loss of the secondary. For example, most people would agree that the loss of braking (a secondary function of a car) is at least, if not more important than the loss of drive (the primary function). Similarly, for a gas compressor, the loss of containment of explosive gas (a secondary function) would normally be considered at least as important as the loss of compression (the primary function).

The consequences of all failure modes – whatever function they cause the loss of – are assessed objectively later in the RCM process.

Should the primary function(s) and secondary functions be labelled as primary or secondary?

The principle of primary and secondary functions is introduced because the primary function(s) is an efficient and clean place to start. The primary function often covers many system components, and often has the highest number of failure modes. However, all functions are analysed so it is not necessary to label functions as primary or secondary. (However, if the user wishes to do so, and this does not take excessive time or cause excessive friction, labelling functions as primary and secondary is not harmful either.) It follows from the above that the group should not argue about whether a function is primary or secondary. If a function is a new and unique requirement it is better just to list it and move on.

2.3 At what level should functions be written?

Just as a system has an asset hierarchy, a system also has a Functional Hierarchy. In other words, Functions can be written at different levels. A (cut down) example is given below for a car/ engine/ crankshaft. As we move down the asset hierarchy, we also have functions written at different levels.

This raises the question; At what level should functions be written? The answer is that they should be written at as high a level as possible, for the following reasons:

1. A greater number of components are covered if the function is written at a higher level, so the overall number of functions recorded in the analysis will be lower (and therefore more manageable).
2. From a user perspective, functions at a higher level are more relevant (a driver of the car is usually more interested in the speed of the car than the load on the crankshaft bearings!)

3. If functions are recorded at too low a level, performance standards get increasingly hard to define.
4. If the asset is broken down to too low a level, then some components may get missed, or their functionality duplicated, which is inefficient and ineffective.

Note however, that sometimes components that are low down in the Asset Hierarchy give additional functionality to the system. In this case, its function is also a function of the system and should be recorded as such. For example, the indicator on a car would be a function of a car, as would the low oil pressure warning light, because these are both distinct and unique requirements (assuming they are not covered by the primary function or any of the higher-level secondary functions listed previously).

The next section describes an efficient approach to functions definition, that starts at a high level and then goes systematically through the asset to determine whether lower-level components give additional functionality.

2.4 How to identify Functions systematically

While it is often argued that the best time to do RCM is at the design stage, the process is usually applied to existing assets. In order to identify a comprehensive list of functions it is beneficial to adopt a systematic approach, to ensure nothing is missed. Such an approach is described in the diagram below.

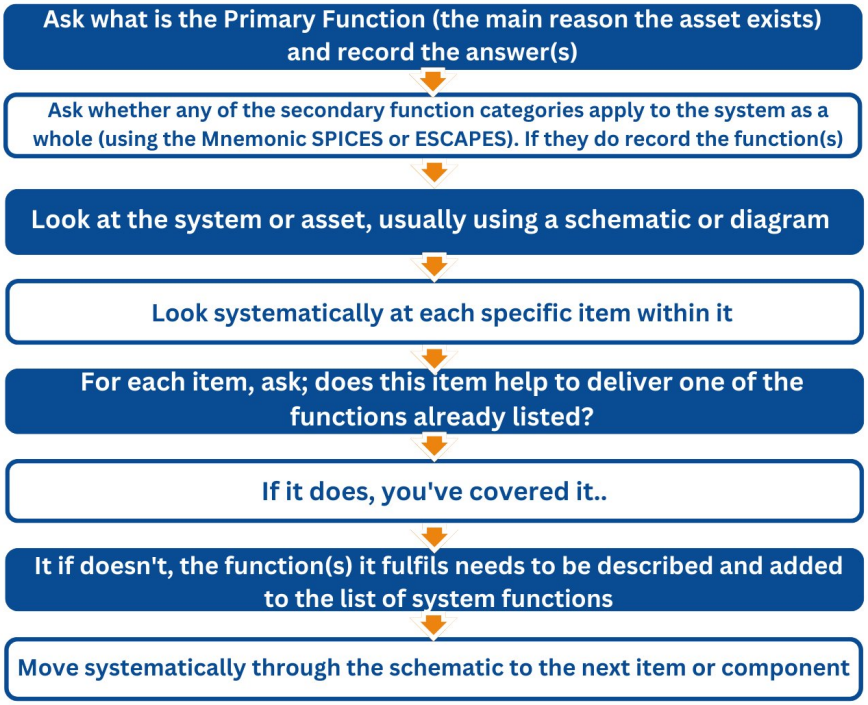

Figure 12: Systemic approach to identifying functions

2.5 The distinction between a functional specification and a design specification.

The Bauhaus Design principle (1925) is still used today and states that "Form follows Function". The form is what the asset is or what is must be, while the function is what it

must do. There is a close relationship, but a subtle and important distinction, between a functional specification and a design specification.

Figure 11: Examples of Form following Function

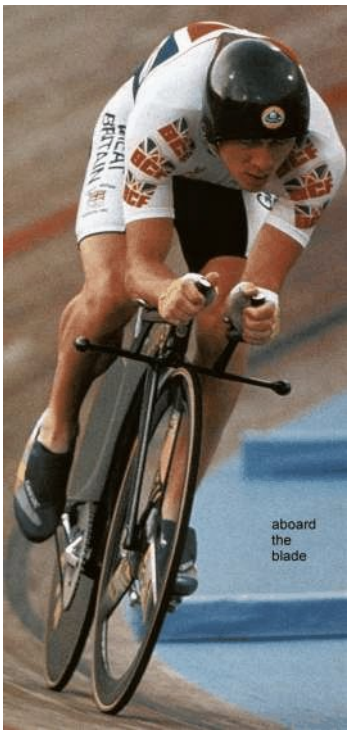

0

When Functions are listed, a functional specification is written. This outlines the expected behaviour of the system under various conditions and describes its interactions with users and with other systems. It's about the "what" and the "why" – what the system should do and why it should do it.

On the other hand, a design specification describes how the asset is constructed to meet the functional specifications. It is about the "how". It describes what the system must be.

While the functional specification is more about the business needs, the design specification is more about the technical solutions to meet those needs. The functional specification ensures that the product meets the needs of the users and the business, while the design specification guides the developers in building the product.

Some RCM analysts incorrectly define the design when writing functions, rather than what the equipment is supposed to do. For example:

Asset	Incorrect Function Statement	Correct function statement
Duty Pump	To provide a pump	To pump water from tank A to pump B at a rate of at least 300 lpm
Standby Pump	To provide a standby pump	To pump water from tank A to pump B at a rate of at least 300 lpm in the event of failure of the duty pump
Warning sign	To provide a warning sign	To warn personnel of the risk of electrocution in the vicinity of the switchgear

2.6 Performance standards

Standard(s) of performance which we want our assets to be able to achieve must be included in the function statements, where possible. This is necessary to provide some clarity about when an asset is working and when it has failed. Sometimes there might be a tolerance, e.g. 40 barg +/- 2 barg; 230V +/- 10%, sometimes there is only a minimum or maximum criterion.

Sometimes we want absolute performance, so a standard does not have to be stated, for example, “To contain the insulating oil” or “To prevent electrical shock”

In other cases a standard cannot be defined, so a subjective standard is used, for example “To look acceptable”.

If a range of performances are required, the function statement should describe the worst case requirement. For example, “To carry up to 5 people at speeds up to 120 kph....”

2.7 Writing functions

When writing functions, we do not have to worry about grouping them under suggested categories, e.g. “control”, “safety” or “protection”.

Functions require a verb and are written in the infinitive – we start with the wording ‘To do....’ (or ‘Not to do....’).

Importance of precision and wording Precision and clear wording are important when writing functions because it is crucial to avoid misinterpretation. Ambiguity can lead to differing interpretations of the requirements, causing confusion among team members and potentially leading to an incorrect decision later on in the process. Clear wording helps ensure that everyone understands the requirements in the same way.

Are Functions ‘set in stone’ for the whole of the analysis?

It is not unusual for adjustments to be made to functions statements as the analysis progresses.

How long should functions development take?

It is not unusual for a good deal of time to be spent discussing, defining and refining the system functions. Sometimes a key word in the function may have great significance and be the subject of heated discussion. If it aids understanding and builds consensus, time spent this way is not wasted time. The Functions definition can take up to a third of the total analysis time, although in the author’s experience, it is usually significantly less than this.

How many functions should a system have?

This depends on the size of the analysis. The smallest analyses may have a handful of functions. If a system is larger, particularly with many protective devices, the number of functions may be significantly above 20. However, if the number approaches 100, something may be wrong (even if it is just that the analysis scope if too broad). A common cause of excessive numbers of functions is writing them at too low a level.

The role of the facilitator when writing functions The role of the RCM facilitator is covered later in the book in greater depth. One point to note regarding functions; the facilitator is not just a 'scribe' who writes down what the group says. Sometimes some interpretation, summarising and careful wording needs to be carried out, particularly when developing functions, and the facilitator should lead this.

2.8 The value and importance of writing functions

When reviewing the functions of an asset is it necessary to consider how the asset works. This leads to improved understanding about the asset by the team members. It also often leads to surprising and valuable discoveries. The importance of this can be illustrated from the following statement taken from the Baker report into the Texas City refinery incident (P3) 'When people lose an appreciation of how their safety systems were intended to work, safety systems and controls can deteriorate, lessons can be forgotten, and hazards and deviations from safe operating procedures can be accepted'

A frequent comments team members make after an RCM analysis is how much they have learnt about the system or asset. This happens even to people who have worked with the asset for many years. Learning opportunities apply throughout the RCM process, but one stage where this happens a lot, is at the functions stage, where the team is asked to describe; how does this component/ sub system work and what do we want it to do?

'Softer' benefits like improved understanding about the asset ultimately translate into 'harder' benefits such as improved reliability, availability and cost.

If functions are well written the Functional Failures are easy to define – this is the next stage of RCM.

2.9 Functions Examples

The following subsections give two Functions examples – the central heating system and the distribution transformer.

2.9.1 Central Heating System

1. To maintain a comfortable indoor temperature of 18-22°C, meeting a temperature control standard of +/- 1°C.
2. To provide consistent heat output to all parts of the building with a total heating output of at least 10 kW.
3. To achieve an overall energy efficiency of at least 90%
4. To prevent the temperature of the water in the boiler primary loop exceeding 90°C.
5. To relieve if the pressure anywhere in the system exceeds 3 bar.

6. To indicate primary and secondary loop temperature and pressure to an accuracy of +/- 1%
7. To provide hot water for domestic use, meeting a hot water output standard of at least 10 liters per minute.
8. To circulate water throughout the system, meeting a minimum flow rate standard of 6 liters per minute.
9. To allow air to be bled from the system
10. To allow the system to be topped up with mains water
11. To prevent corrosion and scaling within the system, meeting an inhibitor concentration standard of at least 10%.
12. To comply with relevant safety regulations and standards, such as EN 12828, EN 14336, and BS 7593.
13. To alarm if the carbon monoxide level, at any point around the boiler exceeds 50 ppm.
14. To vent exhaust gases and condensed fluids safely to the outside environment

2.9.2 Functions Example; Distribution Transformer

1. To step down 50Hz supply from 20kV to 400V at a rate up to 630 kVA on all three phases
2. To transform with an energy efficiency of at least 97%.
3. To contain oil
4. To cool the oil to less than 90 Deg C anywhere in the system
5. To indicate the transformer oil and HV and LV winding temperature locally
6. To trip the transformer if the oil temperature exceeds 105 Deg C (alarm 90 Deg C), or the winding temperature exceeds 130 Deg C (alarm 115 Deg C).
7. To initiate a trip if there is a high flow of oil between the main tank and the expansion tank, or the HV cable boxes and the expansion tank
8. To indicate the oil level in the conservator locally and send a signal to SCADA if the oil level is too low
9. To allow breathing while removing excess moisture and contaminants
10. To relieve excess pressure in the main tank and the cable boxes
11. To provide local indications of transformer model and Asset ID, and signs to warn personnel and members of the pubic of risk of electrocution.
12. To accommodate changes in volume of oil as the temperature changes
13. To provide structural support
14. To operate with a maximum noise level standard of 50 decibels.
15. To allow for easy access to the transformer for maintenance and repair purposes, meeting accessibility standards set by regulatory bodies.
16. To withstand short-circuit currents, meeting a maximum short-circuit current standard of 25 kA.
17. To prevent damage from lightning strikes, meeting lightning protection standards set by regulatory bodies.
18. To provide sufficient clearances around the transformer for safety and maintenance purposes, meeting clearance standards set by regulatory bodies.

19. To comply with relevant safety regulations and standards, such as IEC 60076, EN 50341, and EN 60296.

CHAPTER 3

Functional Failures, Failure Modes and Failure Effects

The 7 master questions of RCM are represented below. This chapter addresses the second, third and fourth of these questions: Functional Failures, Failure Modes and Failure Effects.

1. What do we want the asset to do? – Functions
2. How can it fail? – Functional Failures
3. What causes the Functional Failures? – Failure Modes
4. What happens when a failure occurs? – Failure Effects
5. How much does each failure matter? – Consequences
6. Can we predict or prevent failure and should we be doing so? – Proactive Tasks
7. How should we manage the failure if prediction or prevention is not an option? – Default Tasks

3.1 Functional Failures

The objectives of the failure management programme are defined by the functions and associated performance expectations of the asset. But how are these objectives achieved? The only occurrence that is likely to stop any asset performing to the standard required by its users is some kind of failure. This suggests that maintenance achieves its objectives by adopting a suitable approach to the management of failure. However, before we can apply a suitable blend of failure management tools, we need to identify what failures can occur. RCM does this in two stages:

- Firstly, by identifying what circumstances amount to a failed state. These are the Functional Failures
- Then by asking what events can cause the asset to get into a failed state. These are the Failure Modes.

In 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).

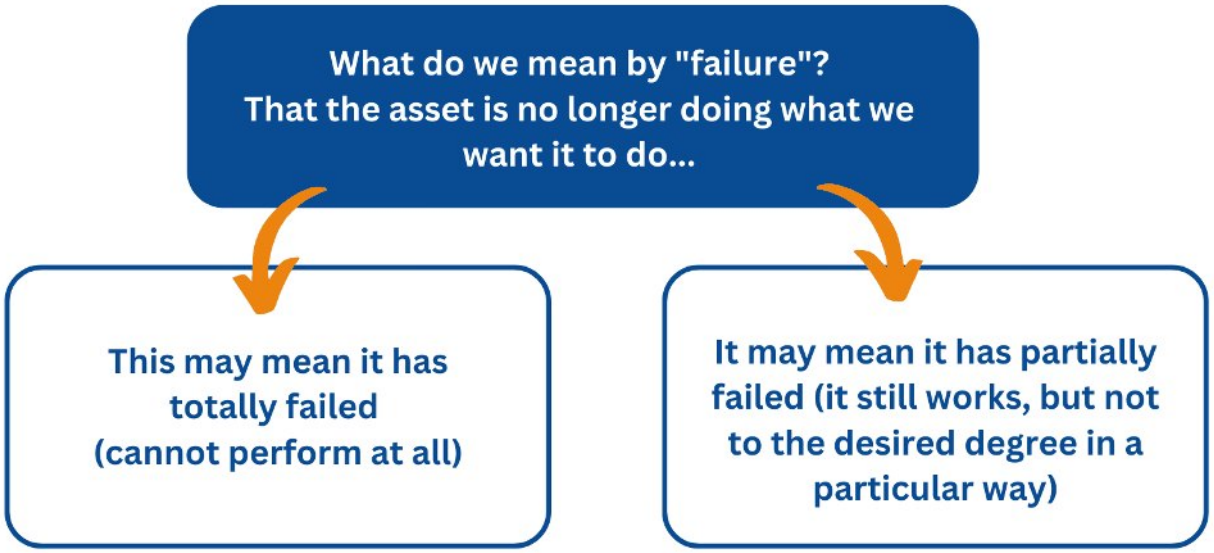

Figure 12: The concept of failure

If the functions are constructed and written well, it will not be difficult to determine the Functional Failures – they are the negatives of the function, and there is a Functional Failure for each performance standard of the function. It should be remembered to make the distinction between total and partial failure.

Function	Functional Failures__
1 To transport up to 4 people from A to B at speeds up to 90 mph	A Unable to transport at all
	B Unable to transport 4 people
	C Unable to transport at 90 mph
2 To stop the car from 30 mph to rest within 120' on the dry road	A To transport up to 4 people from A to B at speeds up to 90 mph
3 To limit discharge of CO2 in the exhaust to less than 170g/km	A To transport up to 4 people from A to B at speeds up to 90 mph

Figure 13: Example Functions and Functional Failures

Note that protective devices can be in a failed state in two ways – operating when they shouldn't and failing to operate when they should – as shown in functional failures fiB and fiB below.

Function	Functional Failures_
1 To prevent any leakage of process fluid from the pump	A Allows leakage of process fluid from the pump
2 To annunciate in the control room if the nitrogen pressure falls below 18 bar	A Fails to annunciate in the control room if the nitrogen pressure falls below 18 bar B Annunciates in the control room when the nitrogen pressure is above 18 bar
3 To indicate nitrogen pressure to within an accuracy of +/- 0.1 bar	A Fails to indicate nitrogen pressure at all B Does not indicates nitrogen pressure to within an accuracy of +/- 0.1 bar
4 To relieve seal pressure >24 bar	A Unable to relieve the seal pressure > 24 bar B Relieves seal pressure below 24 bar

Figure 14: Example Functions and Functional Failures

3.2 Failure Modes

So far, we have specified what we want from the asset (Functions) and the various ways we can be let down (Functional Failures), but we do not yet know how to manage reliability. A key step in the successful management of Reliability is the definition of Failure Modes. This because Reliability is managed at the level of the Failure Mode - Failure Management Tasks manage Failure Modes and it is Failure Modes that the RCM Decision Diagram is applied to.

3.2.1 What are Failure Modes?

Failure Modes are the causes of the Functional Failures.

3.2.2 Why is the correct wording of the Failure Mode important?

If the causes of failure are to be managed appropriately, they need to be clearly understood and correctly defined. If Failure Mode descriptions are vague and ambiguous, either the task may not be properly defined, it may be incorrect, or the connection between the task and the Failure Mode is not apparent, so there is no clear audit trail. For example, if the failure mode is written as 'Pump Fails' there is no clear connection to any of the failure management tasks listed below (even if they are appropriate):

- Check pump bearing for vibration every month using hand-held vibration monitor
- Monitor pump flow by recording the time to fill the tank from A to B every 3 months
- Visual check for leakage of water from the seals

Therefore, if any of the tasks above were proposed for that particular Failure Mode, either the task description, or (more likely), the Failure Mode description would need to be adjusted.

Some examples of appropriate Failure Mode Descriptions with associated Failure Management tasks are given below (note that for these examples, there could be several failure management task options, and some are better than others, but at this stage of the book we will not concern ourselves about which is the best task for each Failure Mode – this will be discussed later in the book):

	Failure Mode	Failure Management Task
1	Pump bearing fails due to normal wear	Check pump bearing for vibration every month using hand-held vibration monitor
2	Pump impeller fails due to normal erosion	Monitor pump flow by recording the time to fill the tank from A to B every 3 months
3	Pump impeller fitted incorrectly	Ensure pump impeller is fitted correctly during maintenance
4	Pump motor fails	No scheduled maintenance
5	Standby pump fails to start	Check standby pump starts when the duty pump trips, every month

3.2.3 Failure Patterns and Failure Modes

We will see later that some failure management tasks only manage certain types of failure pattern. Failure Modes may have different patterns of failure and sometimes the failure pattern may be implied from the failure mode description. For example, an abstraction pump impeller might fail in different ways:

Figure 15: Different patterns of failure for Failure Modes experienced by the same component

Another good example is a car tyre which may fail in different ways:

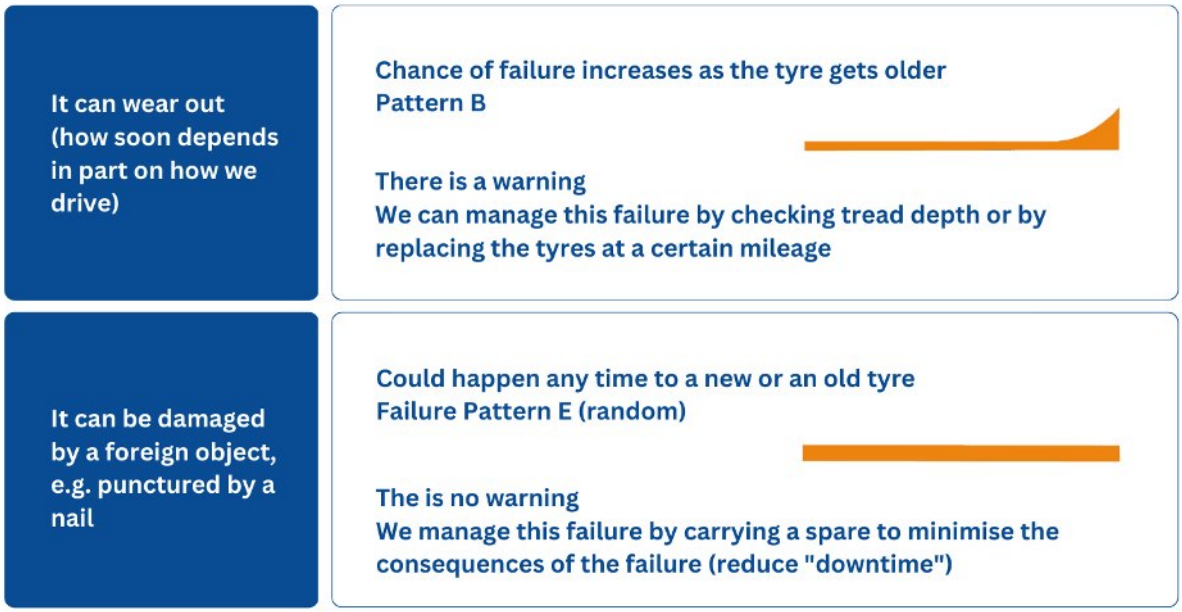

Figure 16: The Failure Pattern may sometimes be implied by the Failure Mode

Do you need to record the failure pattern(s) for every Failure Mode?

The failure pattern implied by the description of the Failure Mode provides important information to decide what should be done to manage the Functional Failure. However, it is not a requirement of RCM to record the failure pattern(s) for every Failure Mode. Some software packages have this option, and it can be useful because it filters out failure management tasks that are not appropriate for certain failure patterns. If you have a software package that does this, it is a useful feature, but it is not an essential requirement to do RCM properly (the appropriate questions related to the failure pattern come up at the relevant stage of the application of the Decision Diagram as we shall see later).

3.2.4 Failure Modes and the Operating Context

To determine a maintenance programme that it appropriate for an asset in its operating context, Failure Modes must be relevant to the asset in its context. For example, the Failure Modes for a duty pump will be different to those for a standby pump – and the failure management policy derived as a result will also be different. Similarly, a transformer in a remote location in an inland hot desert environment may have different failure modes, and different failure management strategies, to a transformer in a in a cold and damp Northern European city centre.

3.2.5 Different types of Failure Modes

Identifying the cause of failure is important because we need to correctly understand the cause to manage the failure properly. In RCM, all causes of failure are considered. This goes back to the Third Insight - asset failure is the product of many things and is driven by 4 forces:

- Deterioration – what the asset is capable of doing deteriorates as time goes by
- Over time, the WANT increases and may even exceed the CAN
- The equipment is incapable of doing what we want it to do from Day one

- Some form of Human Involvement causes the Asset to be incapable of doing what we want it to do.

If any of the four forces causes the Functional Failure, the relevant Failure Modes should be recorded. This raises an important point, that will be made several times throughout the book – RCM is not just about identifying maintenance tasks and their intervals. Rather it is about identifying the most appropriate mixture of failure management tasks (including maintenance, physical design, procedures, training, spares, contingency etc) to manage each reasonably likely Failure Mode. In other words, it is about Reliability, and not just Maintenance. Therefore, Failure modes caused by any of the four 'forces' should be listed if applicable. These are explored further below.

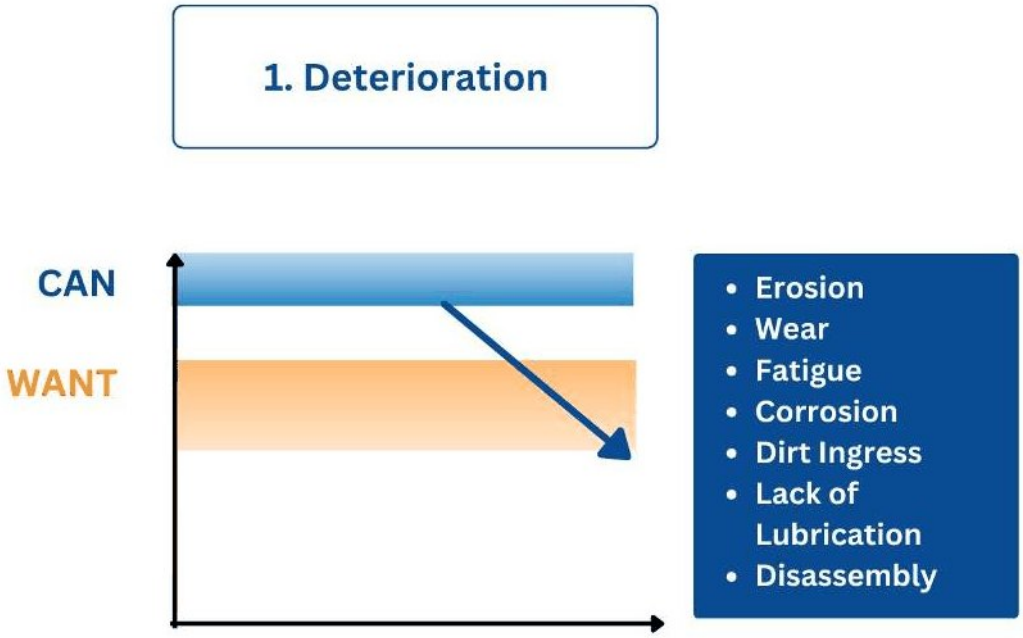

3.2.5.1 'Can' Falling below 'want'

These mechanisms do not need much further explanation. Just before he died, John Moubray explored a fundamental scientific basis for deterioration of systems. He wrote a paper entitled 'Time's Arrow' which went back to the second law of thermodynamics.** The second law of thermodynamics states that:

> Energy spontaneously tends to flow only from being concentrated in one place to

> *becoming diffused and spread out.**

In other words, systems naturally progress towards a more disordered or random state. Maintenance involves putting energy into systems to restore order. Many of the mechanisms that cause disorder are those that cause 'can' to fall below 'want'.

> **The Second Law of Thermodynamics is sometimes referred to as "time's arrow" because it gives a direction to the flow of time. In simple terms, it states that in any closed system, the amount of disorder, or entropy, tends to increase over time.

> Let's use the example of a cup of hot coffee in a cooler room. Over time, the coffee cools down, and the room warms up a little until the temperature is equal. This represents an increase in entropy because the energy is spreading out, becoming more evenly distributed. No matter how long you wait, the coffee won't spontaneously become hot again, and the room won't return to its original cooler temperature. That's because time only goes forward, not backward. This unidirectional nature of time, as explained by the second law, is why it's often referred to as "time's arrow."
>
> So, in a sense, the Second Law of Thermodynamics gives us a way to distinguish the past (when the coffee was hot) from the future (when the coffee is cold). This distinction provides us with an "arrow of time.

https://www.efficientplantmag.com/2004/09/times-arrow/

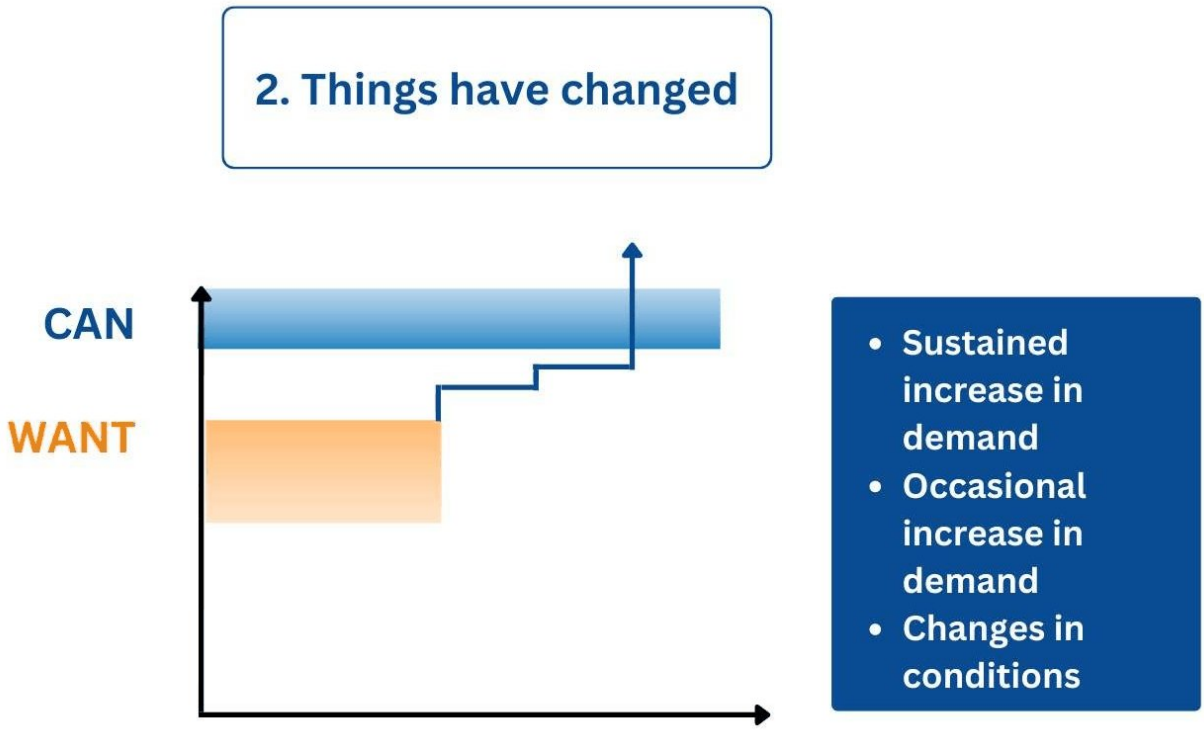

> *F. L. Lambert, "The Second Law of Thermodynamics," March 2003.*

3.2.5.2 'Want' rising above 'can'

This is very common. If a company can sell all that it can make, there will almost inevitably be pressure to increase output. Operations and maintenance personnel frequently complain that assets are operating beyond their original capability. It seems that some plants keep increasing output until some part of the process no longer copes! An important point about this phenomenon is that it is not just about assets as a whole operating (or trying to operate) beyond their capacity. The principle applies at different levels (eg system, sub-system and component), and for different parameters. Such parameters include molecular mass, pH, particle count, cleanliness, levels corrosive or abrasive materials, flow, pressure, temperature etc etc. Whenever something changes in a process, there is often a reason some asset may have too much asked of it in some way. A few examples include:

- A motor driving a gas compressor is overloaded because the gas composition changes as a gas field ages

- A packing machine fails because recycled cardboard behaves differently to the virgin cardboard that was used previously
- A water pump impeller erodes too quickly because the particle count in the water increases beyond the capability of the pump impeller material to deal with it
- Some labels use heat as part of the adhesion process. They require certain thermal characteristics for the container they adhere to. If this changes (perhaps due to a change to a recyclable material for the container) the labels may not adhere properly
- Components in an electrical distribution system become overloaded due to increases in demand resulting from new developments in the vicinity, or because more customers fit air-conditioning
- A new supplier of ink has been found by the purchasing department that is much cheaper. Unfortunately, the ink is not as strong as the old version. In some cases, equipment suffers because it operates at a lower capacity than it should. Examples include:

1. Internal Combustion Engines: These are often designed to operate most efficiently within a certain RPM range. If they are regularly operated at a low RPM, they can experience a build-up of carbon deposits, increased wear due to incomplete combustion, or other issues.
2. Boilers: Boilers can suffer from various issues if not regularly operated at their intended capacity. This can include the build-up of deposits on heat exchange surfaces (which can reduce efficiency and increase the risk of failure) or problems with components designed to manage the by-products of combustion.
3. Power Generators: Power generators, especially diesel generators, that are not run at high enough loads for prolonged periods can suffer from a condition known as "wet stacking." This happens when unburned fuel builds up in the exhaust system due to incomplete combustion at low loads.
4. Air Compressors: If a compressor is oversized for its application and runs at very low loads, it can lead to short cycling which may result in excessive wear on the motor and controls.
5. HVAC Systems: Heating, ventilation, and air conditioning systems can experience various issues if they're consistently used at lower than intended capacities. This might include short cycling, reduced efficiency, or increased wear on components.
6. Pumps: Many types of pumps can experience issues when operated at too low a capacity, which might include overheating, cavitation, or excessive wear
7. A glue dispensing machine clogs, and the glue carbonises because the required flow is too low

3.2.5.3 Want above can right from the beginning

This is more common that it should be! There are many reasons for this. Some reasons and contributing factors from the author's experience are given below (and this list is by no means exhaustive):

- Designers, installers and commissioning personnel are under pressure of time and have the next job to move on to. They face cost pressures to 'get the job done' and move on, so may not have time to properly consider all design options
- Equipment is often selected based on the cheapest initial cost rather than the optimum life-cycle cost
- Companies sub-contract to other organizations who are further removed from the final application (Sub-contractors may themselves subcontract etc etc).
- Communications between departments/ organisations are not always as smooth and comprehensive as they should be.
- Sometimes the original driver/USP for the new product was not properly communicated from the sales/ marketing department, so the product design emphasizes feature X but it is feature Y that the users really want.
- Sometimes user requirements are poorly communicated, or they change (when it is too late), so the 'new' design is not fit-for-purpose when it finally hits the market.
- Usually, the people who buy the equipment are not the people who use it, and often do not ask the people who use it, so the design is sub-optimal.
- Often maintainability and operability are not properly considered at the design stage and are an afterthought.
- Sometimes a sub-optimal design may be specified because it is available/can be made quicker/ is what we did for the last job etc.

- Sometimes space is limited so the equipment has to 'fit in' so the design configuration/operability/maintainability is sub-optimal
- Equipment is made in a low-cost environment but the people making the equipment do not have the correct training
- Equipment is partly made in a low-cost environment and the semi-finished product is shipped around the world (sometimes in several different stages) with all the associated risk that goes with this

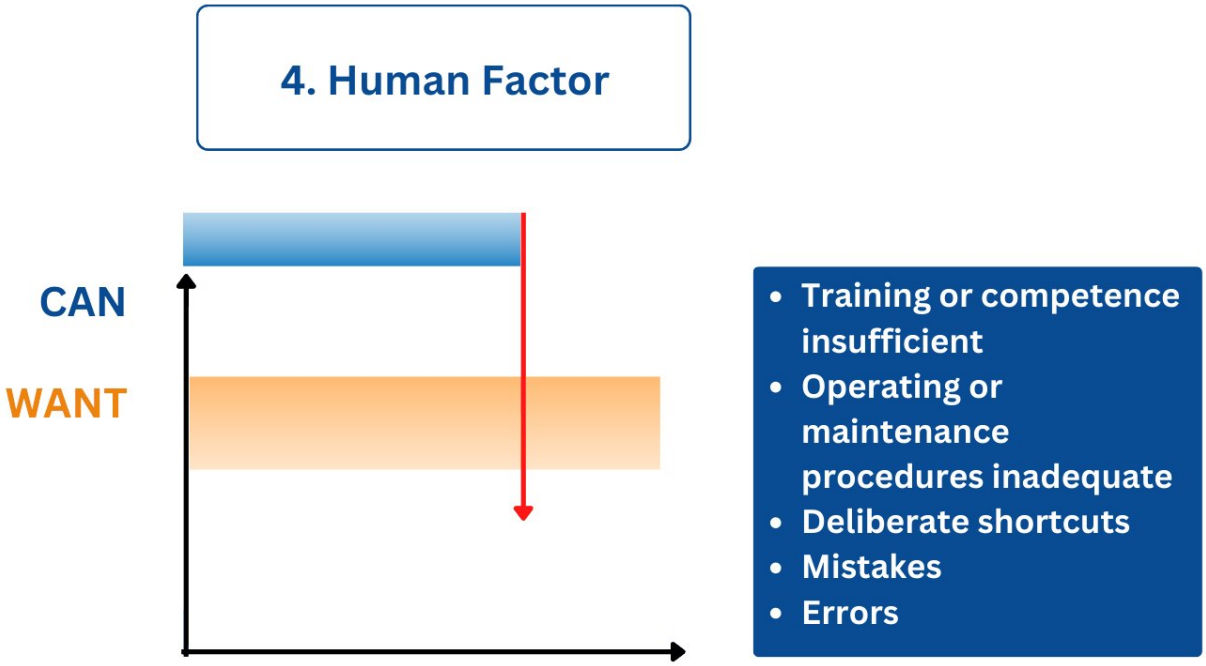

3.2.5.4 Human Error causing 'can' to be below 'want'

James Reason, a renowned British psychologist, is known for his work in the field of human error and complex systems. He developed the 'Swiss Cheese' model of system accidents, helping us understand how accidents occur in industries such as healthcare, nuclear power, and aviation.

Reason classified human errors into two main categories: Active Failures and Latent Conditions.

1. Active Failures: These are unsafe acts committed by people who are in direct contact with the system. They can be categorized into three types:
   - Slips and Lapses (Skill-Based Errors): These are errors in tasks that a person performs often and where the actions are carried out without conscious thought. Slips occur when actions do not go as planned, while lapses occur when planned actions are forgotten.
   - Mistakes (Knowledge-Based and Rule-Based Errors): Mistakes are errors in planning actions or in failing to achieve intended outcomes. Rule-based mistakes occur when a good rule is applied incorrectly or when a bad rule is applied. Knowledge-based mistakes occur when problem-solving processes go wrong.
   - Violations: These are deliberate deviations from recommended procedures or standards. Violations can be routine, where deviation occurs regularly, or they can be situational, where deviation occurs in response to particular circumstances.

2. Latent Conditions: These are indirect failures, often embedded in the system. They arise from decisions made by designers, builders, procedure writers, and top-level management. These decisions may have unintended consequences that can make the system vulnerable to critical failures. They may lie dormant within the system for a long time, only becoming evident when they combine with other factors to lead to an accident.

Remember that James Reason's model of human error is systemic, viewing human errors as consequences rather than causes, indicating that errors are induced by deeper factors within a system. His model suggests that to prevent accidents, it is crucial to manage both active failures and latent conditions effectively.

RCM Case Study

Equipment

Aluminium Recycling Plant

Situation

The process used a simple pump to circulate molten aluminium. This was made from sintered carbon. The pump was unreliable and an RCM study was carried out.

Results

The majority of the failure modes identified during the study were caused by people. These included:

- Damage due to transport within the factory caused by poor practices and inappropriate packaging
- Components damaged in stores caused by poor practices and

The above example raises further points:

- Sometimes there is value looking not just at the equipment itself, but the 'process around the equipment'. If the analysis had just looked at the pumps in isolation, many of the improvement opportunities would not have been identified. However:

o When considering such failures is it important to keep the recommendations sensible and implementable.

o When such failures and considered it is important not to stray too far from the analysis boundaries.

- When considering human error in general, there is almost an infinite number of failure mode possibilities. Therefore, it is critical to ensure that failure modes listed are those that are reasonably likely to occur in the operating context. Doping so requires input from the people that are most familiar with the equipment and its day-to-day operation and maintenance.
- To successfully identify failure caused by human factors, a culture of blame will not work. The analysis is an exercise in identifying problems and recommending improvements to make things better, not a finger-pointing exercise.

3.2.6 Describing Failure Modes

The description of the Failure Mode needs to identify the affected item - system, subsystem or component. It is necessary to identify the cause with sufficient clarity and to a sufficient level of causation such as:

- Turbine blade securing bolt fails due to fatigue
- Gearbox seized due to lack of lubrication

Sometimes a lot of detail is needed and a detailed root cause is appropriate, while sometimes it is appropriate to describe the Failure Mode at a high level and use the term 'failed'. This is often called 'black-boxing'.

The main determinant of the level of detail required in the Failure Mode description is what will be done to manage the Failure Mode – if the task manages the failure at a high level the word 'failed' may be appropriate, if a particular root cause is managed by the task, that root cause should be included in the Failure Mode description.

Identifying the Correct component in the failure mode description

Part of the correct wording of the failure mode is to give the component its correct name. Rather than just 'Pressure gauge drifts' or 'Valve sticks', if there are many gauges or valves on the system, it is better to give the Instrument or valve its proper name (or even its number) in the failure mode description to avoid confusion later on.

3.2.7 Level of detail when writing Failure Modes

In principle it is straightforward to identify Failure Modes – they are the causes of the Functional Failures. However, this is an area where things can go badly wrong. One common problem is to identify the level at which the Failure Mode should be written.

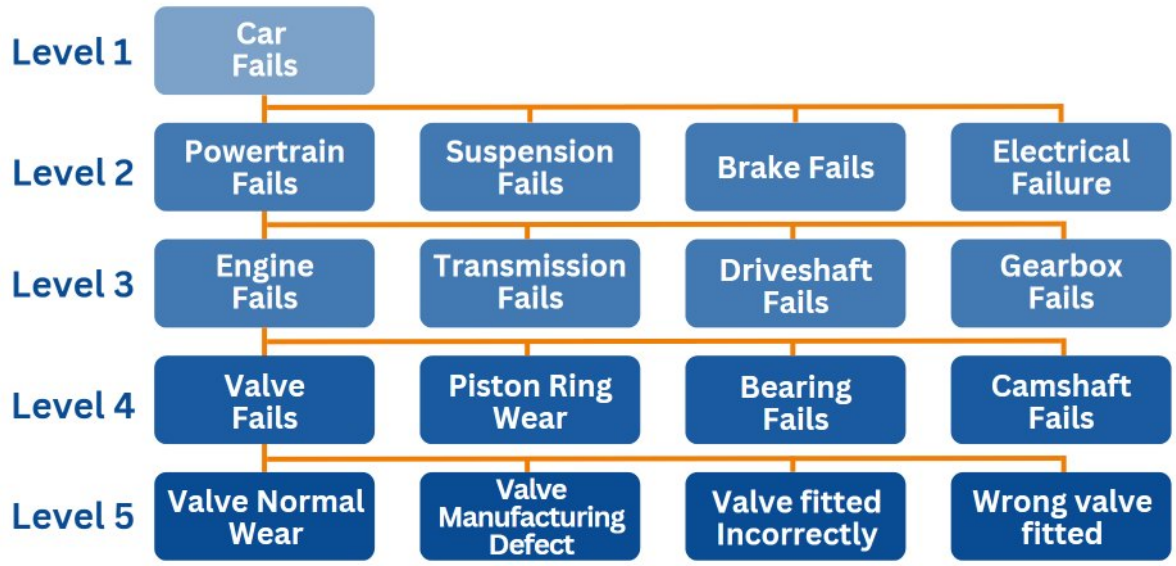

The diagram above shows failure modes for a car at different levels. The top level is a single high level Failure Mode at the level of the asset. While this could be recorded quickly and easily, it would not be much use for developing a robust failure management task as it is too general and vague. If we tried to develop a failure management task based on this state-

ment, failures might be missed or may not be addressed appropriately, leading to increased downtime, risk, and costs.

Therefore, we might want to go to level 2 by asking 'why?' and listing failure modes for systems. There might be 10-20 such statements for a car. Once again, these may be too vague to come up with robust failure management tasks, so we might ask 'why?' again, and move to level 3 and list sub-system failure modes, or we might keep going, and ask 'why?' many more times.

As we move from higher levels of detail to lower levels, the failure modes become more specific and focused on individual components and sub-components and the causes become more specific. These lower-level statements are much more useful for defining robust and auditable failure management tasks. For example, for the failure mode statement Wrong valve)tted we can have a training or procedure recommendation to make ensure the correct valve is fitted, and for the failure mode Valve manufacturing defect we can make sure the valve is manufactured correctly.

If this exercise was carried out for a car, a single failure mode at level 1 could become several hundred at level 5.

Clearly it is necessary to go into more detail than level one, and identify failure modes in more detail, but can we go too far? Coming back to the example above, if the probability of a valve manufacturing defect is vanishingly small, are we going into unnecessary detail? And what about other failures that have low consequences, such as a door handle, or a headrest - do we need to list large numbers of failure modes for these in great detail? Might an analysis containing too many such failure modes be unmanageable and take too much time to be worthwhile?

Fortunately, we do not need to ask why 5 (or any other specified) times; in RCM we ask why an appropriate number of times.

Failure modes should be recorded in sufficient detail to be able to identify a suitable failure management policy. This means that in some cases we dig deeper and go into more detail, in some cases we do not.

So how do we decide the level of detail to go into?

The extent to which we dig deeper depends on:

- The probability of the failure; we may go into more detail of the failure is more common
- The consequences of the failure; we may go into more detail of the failure has greater consequences
- And we should always consider whether we have enough detail to manage the failure mode appropriately

If these guidelines are followed, a comprehensive and effective study can be carried out efficiently.

Should failure mode descriptions always include the phrase 'due to'?

In general terms it is good practice to identify the root cause of failure to a reasonable level of causation. Decisions should deal with causes rather than effects. A lower level of causation is often expressed via the wording 'due to'. For example:

- Valve sticks due to corrosion
- Bearing seizes due to normal wear
- Care tyre fails due to impact with nail in the road

This is also good practice to identify and confirm failures that occur just because of normal wear and tear, rather than something causing the failure that should not have done, because the failure management task may be different. For example:

Failure Mode	Appropriate failure management task
Bearing seizes due to normal wear	Monitor bearing for vibration
Bearing seizes due to incorrect fitting	Fit bearing correctly
Bearing seizes due to lack of lubrication	Lubricate bearing

Some users try to use the words 'due to' for every failure mode statement. The author's view is that while this is good practice in general it is going to far to recommend universally because:

- Sometimes a number of failure causes may be grouped together at a high level (eg electronics card fails)
- Sometimes trying to include a form of words results in clumsy wording

RCM Consulting Case Study

Equipment

Range of equipment in Power System

Situation

The client (a power company in South America that managed fleets of assets) had been badly advised by on the application of RCM by Management Consultants.

For each RCM analysis the consultants had told the client to:

- List all failure modes that were possible
- Ask why 5 times for every failure mode

Results

Massive spreadsheets were generated with excessive repetition of root causes, many hundreds of rows and a large number of columns. These were extremely

difficult to implement, to update and to template to similar assets.

Lesson

Guidance about how to record failure modes is subtle but extremely important for

a successful analysis:

- Failure modes should be listed that are credible (reasonably likely) in the context in question
- The root cause of failure should be identified in sufficient detail to make a sensible decision to manage the failure mode

3.2.8 Failure Mechanisms, Failure Causes and Failure Modes

A Failure Mode is defined in the SAE JA1011 standard as:

A single event that causes a functional failure

An RCM process complies with SAE JA1011 if it has just one column/ box for Failure

Modes, in which the causes of the functional failure are written to an appropriate level. This is what the Decision Diagram is applied to.

Some users specify (and define) variations on this concept such as the Failure Cause and the Failure Mechanism (which we will not get into here!). Some FMEAs have two or even three columns for two or more of these variations.

For the purpose of failure management the SAE JA1011 approach is simple and elegant. The cause of the failure is identified in sufficient detail to manage appropriately and that is it. It is not necessary to define and specify other variations.

That having been said, if a process specifies these additional variations it is not wrong provided:

- The variations are clearly defined and consistently applied
- Additional time filling in the extra boxes is not excessive
- The decision is made to the root cause at a level that allows a suitable failure management policy to be defined (as per the SAE RCM Standard).

(It is also recognised that for the purpose of the collection of failure data, and data sorting and manipulation, additional variations can be useful.)

3.2.9 Sources of Failure Mode information

There are many sources of Failure Mode information as shown in the diagram below. However, the best source of information is often people's knowledge and experience, which resides in people's heads.

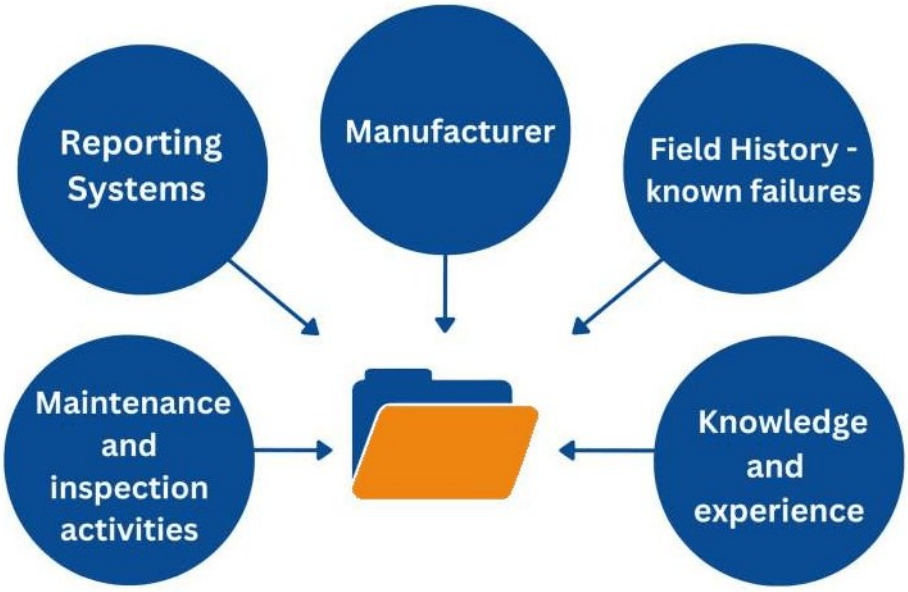

Figure 18: Sources of failure data

3.2.10 Should all conceivable Failure Modes be recorded?

When recording Failure Modes it is necessary to record those that are reasonably likely. It is not cost effective to list all conceivable Failure Modes regardless of likelihood. Clearly, it is important to list those failures that have happened in the past. However, it is also important to look ahead and list Failure Modes that have not yet occurred but which are considered reasonably likely – perhaps because the asset is ageing, or because failures have happened on similar assets in a similar context elsewhere.

Note that failures that are less likely but have greater consequences should also be listed. It is also important to record failures that are presently being prevented by present maintenance activities – in this way RCM can be used to validate or challenge the existing maintenance programme.

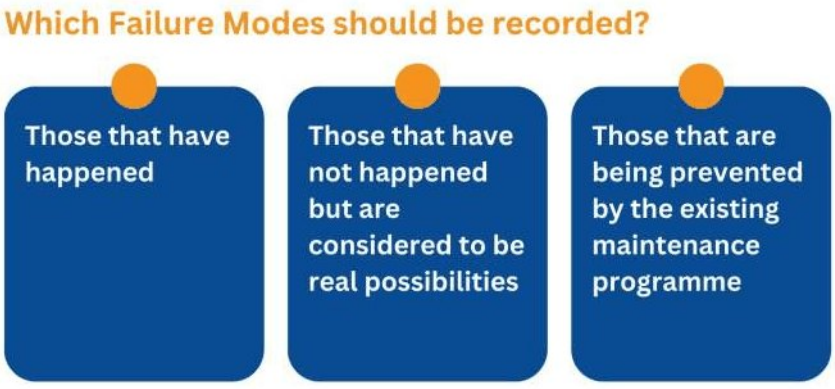

3.2.11 How to Use RCM to review Manufacturers recommendations (or the existing maintenance programme)

Manufacturer's maintenance recommendations are often sub-optimal. The reasons for this are discussed later in the book. RCM can be used to challenge, review, validate and improve manufacturers maintenance recommendations.

The key to doing this is to recognise that all maintenance tasks manage failure modes. For each manufacturer recommended task, the failure mode the task addresses can be inferred by asking why? Or For what reason? For example:

Equipment	Proposed Task	Inferred Failure Mode
Reciprocating Compressor	Visually inspect valve	Valve wear
Filter	Replace filter	Filter clogs
Pump	Listen for noise	Bearing wears
Transformer	Check for leaks from tank	Tank corrodes
High temperature transmitter	Calibrate transmitter	Transmitter drifts

The failure effects can then be described and the Decision Diagram applied to the failure mode to determine the optimum task to manage the failure mode, and the appropriate interval for that task.

The same approach can be used to challenge and review the existing maintenance programme.

The author has carried out many successful projects where improvements have been made to manufacturers maintenance recommendations or the existing maintenance programme by adjustments to the task and/or the interval.

(In some cases, it is difficult to determine exactly what failure mode the recommended task is addressing because the task is too vague and general.)

3.2.12 Criticality Scoring of Failure Modes

Criticality is often expressed as a combination of probability/ likelihood of failure and severity/ consequence. It can be identified for assets as a whole, and also for Failure Modes of assets. A matrix, such as the one shown below, is an example of a criticality matrix.

Frequency / consequence	1 Very unlikely	2 Remote	3 Occasional	4 Probable	5 Frequent
Catastrophic					
Critical					
Major					
Minor					

Matrices can have varying numbers of 'boxes' – many are 5 x 5 but some are 3 x 3 or 7 x 7. Categories of consequence used include safety, the environment, cost, availability, reputation, legal etc. Some matrices have quantitative definitions (ie numbers and figures), others have more qualitative definitions, many have both.

There is no single, universal Criticality (Risk) Matrix. A good criticality matrix is one that:

- ranks the failure modes in order of criticality.
- properly reflects the impact of the failure mode on the business.

Why Score Criticality of failure modes?

Criticality allows prioritization of implementation. Failure Management tasks dealing with Failure Modes that are high criticality should be implemented first, those with lower criticalities can be implemented afterwards. Priority of implementation is particularly useful for planning the implementation of one-off changes/ redesigns.

Some users score criticality unmanaged (ie without any failure management task applied) and managed (ie with a failure management task applied).

> Scoring both managed and unmanaged criticality in FMECA provides several benefits:
>
> 1. Risk Identification and Prioritization: By considering both managed and unmanaged scenarios, you can better understand the overall risk profile of your system. This helps in prioritizing the risk mitigation strategies and in directing resources where they are most needed.
> 2. Effectiveness of Controls: By comparing managed and unmanaged criticality scores, you can gauge the effectiveness of your current mitigation strategies. If the managed criticality score is significantly lower than the unmanaged score, it suggests that your controls are effective.
> 3. Risk Mitigation Strategy Planning: If the managed criticality score is still high, it indicates that further measures are needed to control the risk. This allows you to plan and implement further risk mitigation strategies.
> 4. Understanding Worst-case Scenarios: Assessing unmanaged criticality allows you to understand the worst-case scenario, should all controls fail. This can be useful for contingency planning.
> 5. Compliance: Demonstrating that you've considered both managed and unmanaged scenarios may be a requirement for compliance with certain industry standards or regulations.
> 6. Continuous Improvement: Regularly updating the FMECA to account for changes in both managed and unmanaged criticality can lead to continuous improvement in system design and risk management practices.

FMEA versus FMECA

A Failure Modes and Effects Analysis (FMEA) does not include criticality of failure modes. A Failure Modes, Effects and Criticality Analysis does.

Does RCM include criticality of failure modes?

RCM includes a FMEA. SAE standards JA1011/ 1012 do not specify criticality of failure modes as a requirement for an analysis to be an RCM analysis. Therefore, an RCM analysis can be carried out without scoring criticality of each failure mode.

The reason RCM does not include Failure Mode criticality is that all the elements of criticality (probability of failure, consequences of failure (in all its forms)) are covered as part of the application of the RCM Decision Diagram.

If I do a FMECA as part of my RCM analysis am I still doing RCM?

Providing all the requirements of the SAE standard are complied with, an analysis is an RCM analysis. If Failure Mode criticality is scored as well, the user has gone beyond the requirements but the analysis is still an RCM analysis.

Risk Priority Number (RPN)

> RPN is an extension of criticality. It is a numerical value based on three main factors: severity, occurrence, and detection.
>
> 1. Severity: This represents the seriousness of the failure's effect. Higher values indicate more severe consequences.
> 2. Occurrence: This reflects the frequency with which the failure is expected to occur. A higher value means the failure is more likely to occur.
> 3. Detection: This refers to the likelihood of the failure being detected before it causes a problem. A higher value indicates that detection is less likely.
>
> The RPN is usually calculated by multiplying the values of these three factors. For example, if each of these factors has a score of 1 (good) to 5 (bad), the RPN score ranges from 1 to 125.

3.2.13 'A Silver Shotgun'

Some Management textbooks talk about a 'Silver Bullet'. "A Silver Bullet" is a metaphor

> that signifies a straightforward, instantaneous, and effective solution to a complex problem. It's seen as remarkably potent, slicing through intricate difficulties with ease

and minimal effort.

> Unfortunately, in most cases (in the author's opinion), the idea of a single, easy but potent solution to solve complex and difficult challenges, is just wishful thinking. In any case, if there ever were any silver bullets, they have surely all been used up by now! What is perhaps more appropriate is the notion of the 'silver shotgun'. Rather than a single improvement that makes all the difference, numerous small improvements add

together to make a big difference.

> In the context of RCM and physical assets, these small improvements are better management of individual failure modes. This principle can be illustrated by considering the range of different approaches that would be used to manage the failure modes of the motorcycle shown below. They include maintenance, procedures, training, contingencies, and spares. Put together they ensure that the motorbike is safe, efficient

and reliable.

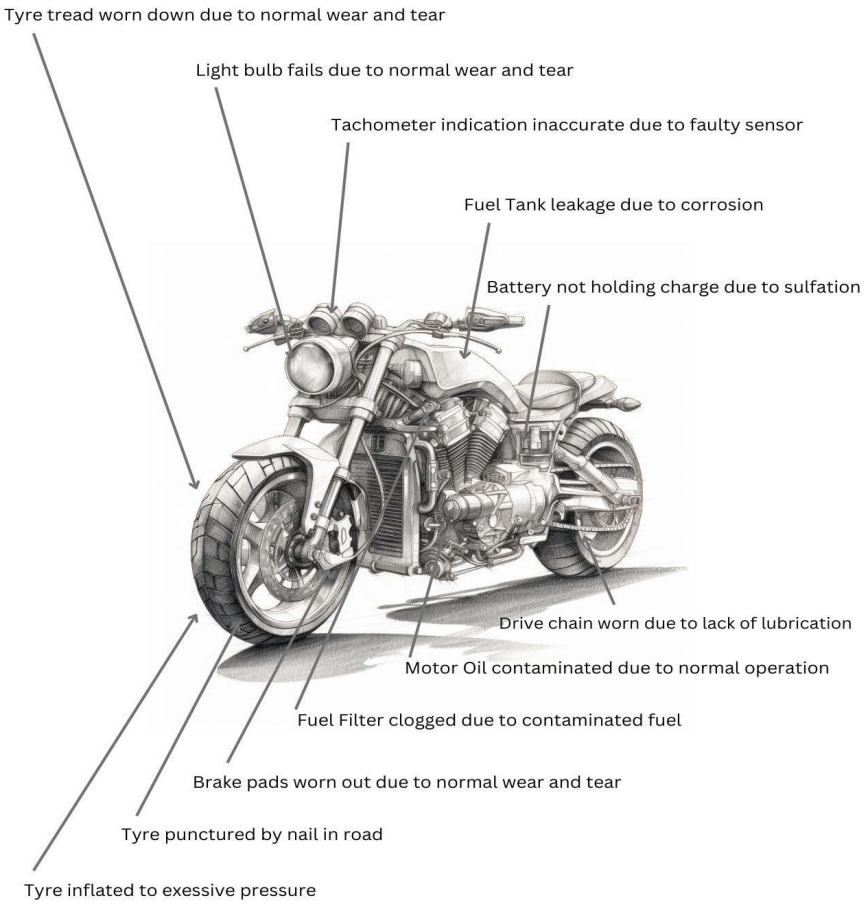

3.3 Failure Effects

Having identified the Failure Modes, we need to understand and document what happens when the failures occur. These are the Failure Effects. We write Failure Effects to help us determine how much the failure matters (its consequences). This is because those failures with major consequences may deserve a lot of attention while those with minor consequences may merit less attention or none at all.

For each Failure Mode we record the Failure Effect. These describe what happens when each failure occurs, in the order that it occurs, to its natural conclusion as if no planned maintenance or inspection was being used to prevent or avoid it.

Where relevant, the following information needs to be described in the Failure Effect:

- any evidence of the failure
- safety / environmental impacts
- effect on operations / process
- secondary damage (if any)
- estimated downtime or time to repair

What happens before the failure mode occurs

Failure effects normally start from the failure mode itself and describe what happens afterwards. It is not a requirement of RCM for warnings before the failure occurs to be recorded. However, some users also describe what happens before the Failure Mode occurs, to encourage the team to consider potential failure conditions that might allow condition-based tasks

to be identified. Since condition-based maintenance is usually the best form of maintenance, it is good practice to do this (provided the ultimate failure management task is not assumed at this stage). Examples of Failure Effects including warnings are shown below:

Failure Mode	Failure Effect	
Crude oil pump bearing fails due to normal wear		Damage to the bearing, loss of lubrication, increased friction, heat
		generation and vibration. Alarm on high bearing temperature and
		vibration leading to a trip. Risk of some damage to the pump mechanism
		due to increased strain on other components. Pump shut down leads to
		shut down of crude oil supply to the process. Time to repair up to 8 hours

Failure Mode	Failure Effect	
Transformer casing fails due to corrosion		Visible signs of corrosion. Leakage of insulating oil and increased risk of
		contamination. Contamination will cause deterioration of internal
		components of the transformer to accelerate. Low oil level can cause
		increased heating of the transformer internals components resulting in an
		alarm and trip on low oil level or high winding temperature. Loss of power
		distribution; potential for environmental hazards due to oil leakage and
		fire risk. Time to repair up to 4 hours

Local, next level and end effects

Some FMEAs have three different 'boxes' for Failure Effects – Local effects, next level effects and end effects. For example:

Failure Mode	Local Effect		Next Level Effect		End Effect	
Crude oil pump bearing fails due to normal wear		Damage to the		Risk of some damage		Pump shut down
		bearing, loss of		to the pump		leads to a shutdown of
		lubrication, increased		mechanism due to		crude oil supply to the

Failure Mode	Local Effect		Next Level Effect		End Effect	
		friction, heat		increased strain on		process. Time to
		generation and		other compon-ents.		repair up to 8 hours

vibration. Alarm on high bearing temperature and vibration leading to a trip.

Failure Mode	Local Effect	Next Level Effect		End Effect
Transformer casing fails due to corrosion	Visible signs of corrosion. Leakage of insulating oil and increased risk of contamina-tion		Contamina-tion will	Loss of power distribution; potential for environmental hazards due to oil leakage and fire risk. Time to repair up to 4 hours
			cause de-terioration of	
			internal compon-ents	
			of the transformer to	
			accelerate. Low oil	
			level can cause	
			increased heating of	
			the trans-former	
			internals compon-ents	
			resulting in an alarm	
			and trip on low oil	
			level or high wind-ing	
			temperat-ure.	

It is not a requirement of RCM to have three boxes like this and most analyses have just one box for Failure Effects. However, important points are that, even if one box is used, Failure

Effects should keep going to the natural conclusion, and should include local, next level and end effects, as shown below:

Failure Mode	Failure Effect	
Crude oil pump bearing fails due to normal wear		Damage to the bearing, loss of lubrication, increased friction, heat
		generation and vibration. Alarm on high bearing temperature and
		vibration leading to a trip. Risk of some damage to the pump mechanism
		due to increased strain on other components. Pump shut down leads to
		shut down of crude oil supply to the process. Time to repair up to 8 hours

Failure Mode	Failure Effect	
Transformer casing fails due to corrosion		Visible signs of corrosion. Leakage of insulating oil and increased risk of
		contamination. Contamination will cause deterioration of internal
		components of the transformer to accelerate. Low oil level can cause
		increased heating of the transformer internals components resulting in an
		alarm and trip on low oil level or high winding temperature. Loss of powe
		distribution; potential for environmental hazards due to oil leakage and
		fire risk. Time to repair up to 4 hours

Key words in failure effects

Some RCM consultants recommend the use of key phrases in the Failure Effects to 'guide' the decisions that follow. Some of these key phrases are as follows:

Phrase	When is it used	What does it (usually) signify
'This only matters if'	If another failure is needed before the failure mode matters or we find out about it	A hidden failure
'Over a period of time'	If something happens progressively leading to a	Some form of predictive or preventive maintenance task

	failure.	may be technically feasible
'Suddenly'	If something happens without warning	Some form of predictive or preventive maintenance task is not technically feasible so failure-finding or redesign may be appropriate

Some examples are as follows:

Failure Mode	Failure Effect	
Transformer casing fails due to corrosion		Over a period of time, visible signs of corrosion develop. Leakage of
		insulating oil and increased risk of contamination. Contamination will
		cause deterioration of internal components of the transformer to
		accelerate. Low oil level can cause increased heating of the transformer
		internals components resulting in an alarm and trip on low oil level or
		high winding temperature. Loss of power distribution; potential for
		environmental hazards due to oil leakage and fire risk. Time to repair up
		to 4 hours

Failure Mode	Failure Effect	
Crude oil pump bearing fails due to incorrect fitting		Suddenly the bearing is fitted incorrectly. This results in damage to the
		bearing, loss of lubrication, increased friction, heat generation and
		vibration. Alarm on high bearing temperature and vibration leading to a
		trip. Risk of some damage to the pump mechanism due to increased
		crude oil supply to the process. Time to repair up to 8 hours

Failure Mode	Failure Effect	
Transformer casing fails due to corrosion		Over a period of time, visible signs of corrosion develop. Leakage of
		insulating oil and increased risk of contamination. Contamination will
		cause deterioration of internal components of the transformer to
		accelerate. Low oil level can cause increased heating of the transformer
		internals components resulting in an alarm and trip on low oil level or

Failure Mode	Failure Effect	
		high winding temperature. Loss of power distribution; potential for
		environmental hazards due to oil leakage and fire risk. Time to repair up
		to 4 hours

Failure Mode	Failure Effect	
Standby lube oil pump motor insulation fails due to normal deterioration		This only matters if the duty pump fails. The standby pump will not cut in.
		Lube oil pressure will continue to fall, a low pressure alarm will sound in
		the control room and the compressor will trip on low lube oil pressure.
		Compressor shut down leads to loss of gas supply to the process. Time
		to repair up to 2 hours

Once again, it is not a requirement of RCM for such phrases to be used. However, they can be useful in the right places, and if they fit nicely, and work well, then it is good practice to use such phrases. The author's view is that it is going too far to insist on certain forms of words in Failure Effects. This is because:

- sometimes attempts to use the given phrases can result in a Failure Effect description that does not read well.
- in other cases, more than one of these phrases may be relevant, leading to unnecessary confusion and complication about what the rules are.
- Unnecessary repetition of the failure mode is sometimes required in the failure effect with the sole purpose of incorporating the correct phrase.

How to handle Protective Devices in Failure Effects

When an FMEA is written each Failure Mode is written one at a time, assuming everything else is working. If we take the example of a pumping system shown below:

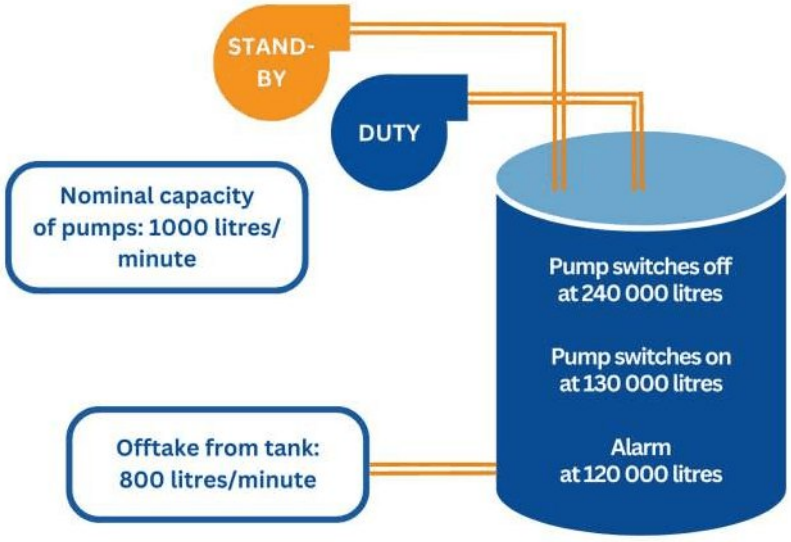

One of the Functions, Functional Failures and Failure Modes and Failure Effects of the duty pump is shown below.

Function	Functional Failure	Failure Mode	Failure Effect
To deliver at least 800 litres/ minute to the process	Fails to deliver at all	Bearing seizes due to normal wear	Pumping stops, level falls, alarm sounds in control room at 120 000 litres. Standby cuts in. Time to replace bearing 4 hours.

Note that the statements:

- Alarm sounds in control room at 120 000 litres.
- Standby cuts in.

Assume that these protective devices work.

These protective devices have separate Functions, Functional Failures, Failure Modes and Effects elsewhere in the analysis. It is at this stage that the failures for these devices are considered and addressed.

Function	Functional Failure	Failure Mode	Failure Effect
To deliver at least 800 litres/ minute to the process in the event of failure of the duty pump	Fails to deliver at least 800 litres/ minute to the process in the event of failure of the duty pump	Standby pump fails to start	This only matters if the duty pump fails. Pumping stops, level falls, alarm sounds in control room at 120 000 litres. Standby cuts in. Time to replace the pump 3 hours.

Function	Functional Failure	Failure Mode	Failure Effect
To sound an alarm in the control room if the level in the tank falls to 120 000 litres	Fails to sound an alarm in the control room if the level in the tank falls to 120 000 litres	Low level alarm fails	If the level falls to 120 000 litres due to a failure elsewhere in the system (such as a pump or a switch) there would be no alarm in the control room. The tank would run dray and production would stop Time to repair the alarm 6 hours.

Repair time

The repair time is the Operations perspective – the time from the equipment stopping until it starts again. Typically, this might comprise of the following steps

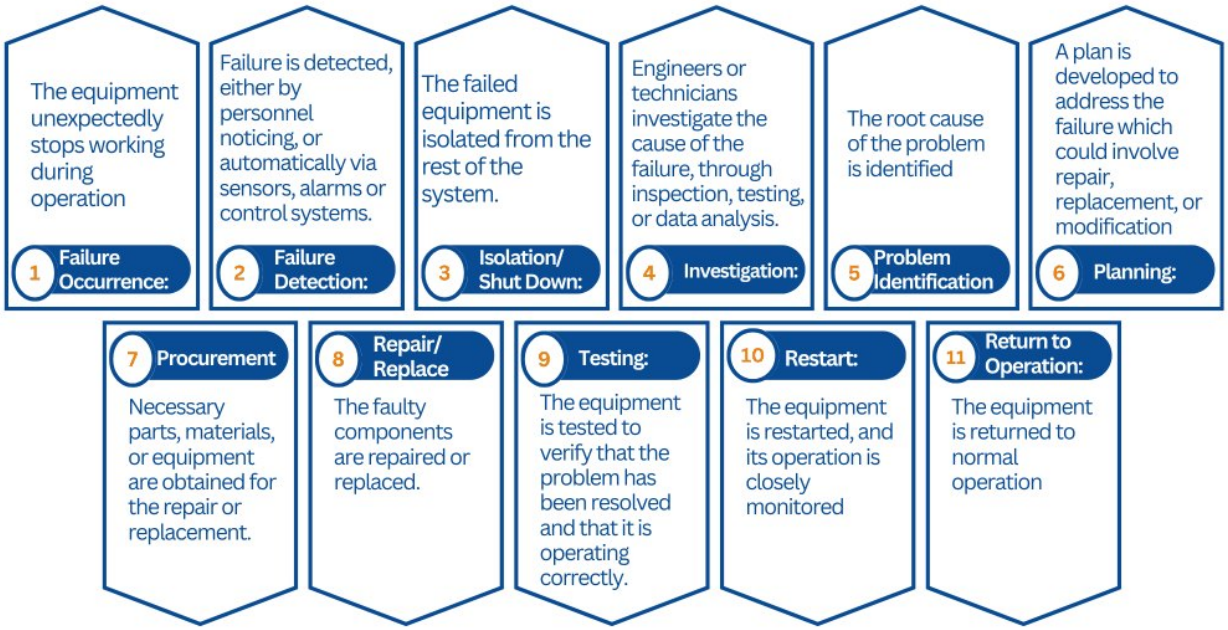

The 'typical worst case'

An important guideline when describing Failure Effects and recording the repair time, is that the 'typical worst case' should be described. The absolute worst case/ perfect storm scenario could result in an overly conservative maintenance programme. But if the average scenario (including Mean Time to Repair) is recorded, the maintenance programme risks being wrong half the time!

Spares and Failure Effects

A big contributor to the repair time can be the time to procure a spare. If a spare is not available a repair time could be months rather than hours. The RCM process is well suited for identifying optimum spares requirements logically and systematically and it is at this stage of the process where the spareshold policy starts to have an impact. An interesting angle on this is that spares and maintenance are not independent. For example:

- if a failure can be predicted with a warning that is less than the leadtime, then the spare can be ordered as soon as the warning is detected, and does not been to be held in stores.
- If a spare is held, the consequences of failure are reduced and certain maintenance tasks may not be worth doing

Zero base and failure effects

Failure effects must be 'zero-based'. In other words, they should describe what would happen as if nothing is being done to prevent or predict the failure. Decisions about failure management come later. Note however, that when it comes to spares, the Failure Effects (including the repair time) should be described with the existing spareshold policy in place (in other words, if the policy is to hold one spare, the Failure Effect should be described as if the spare is available).

The importance of Failure Effects

Failure Effects provide justification for the decisions that are made in the later stages of RCM during the application in the Decision Diagram. For example:

- If the failure consequence is assessed as a safety consequence, this should be clear from the information given in the Failure Effect.
- If the failure consequence is assessed as a hidden consequence, this should be clear from the information given in the Failure Effect.
- If the task is worth doing this should be clear from the information given in the Failure Effect
- If the task is not worth doing this should be clear from the information given in the Failure Effect -

Some examples are given below:

Function	Functional Failure	Failure Mode	Failure Effect	What can be read into the Failure Effect
To sound an alarm if smoke levels in the building exceeds 3%/m of obscuration	Fails to warn residents of smoke if smoke levels in the building exceeds 3%/m of obscuration	Smoke detector fails	This only matters if there is a fire. Workers will not be warned and there is an increased risk they may be killed or injured as a result	This is a hidden failure. The consequence of the multiple failure is safety
To contain hydraulic fluid under pressure	Fails to contain hydraulic fluid under pressure	Flexible hose deteriorates	The hose splits and will spray hydraulic fluid under pressure in the surrounding area. The hose may break completely and fly around violently with risk to personnel in the area.	This is a safety consequence
To control temperature in the oven	Fails to control temperature in the oven	Control thermostat fails closed	Oven would get too hot. Food would be burnt. Loafs would be scrapped.	This is not a safety consequence (because otherwise it would be clear from the effects)

Another interesting aspect regarding Failure Effects is that in addition to providing justification for doing a maintenance task, they provide justification for not doing a maintenance task. An example from the author's experience was an analysis of a seal oil system in a hydrocracker recycle gas compressor at an oil refinery.

RCM Case Study

Equipment

Oil Refinery hydrocracker recycle gas compressor seal oil system

Situation

The primary function of the seal oil system is to provide a tight seal in the area where the compressor's shaft exits the casing. This is crucial in preventing the escape of process gases (which can be hazardous or environmentally harmful) and the ingress of atmospheric air.

This was therefore a critical function with massive safety consequences of failure. The process is also extremely valuable. The system had duty and stand-by pumps with protective devices to warn of pump failure. The system also included accumulators and stand-by accumulators in case of (double) pump failure and for emergency oil supply. The accumulators had protective devices to monitor crucial parameters (such as pressure and level). Therefore, if the accumulators were to fail there would be a warning. Most sensors were smart devices with a high level of self-checking, and redundant devices were fitted. For most of the failure modes of the instruments, the system would identify that they had failed and raise an alarm. The system was fitted with

a high level of protection and redundancy.

Results

For a high proportion of failure modes, the failures were managed through the

design of the system.

This was reflected in the RCM analysis which proposed No Scheduled Maintenance for a large proportion of Failure Modes. (This recommendation wasn't universal - some maintenance was proposed to check protective devices

and to reduce overall system risk).

Lesson

It is arguably a bigger decision to decide not to do a maintenance task to manage a particular failure mode. If this decision is taken because the task is not worth doing, the justification should be clear from the Failure Effect description.

How to do Failure Modes and Effects in the RCM Meetings

Failure Modes and effects are inextricably linked - an effect logically follows a cause. For this reason, it is most efficient to record the Failure Mode, and then record the Failure Effect straight away at the same time.

The difference between Failure Effect and Failure Consequence

> Understanding the facts before making a judgment is a crucial practice in all areas of life, for the following reasons:
>
> 1. Accuracy: A judgment based on solid facts is more likely to be accurate and reliable. Without all the facts, you may make decisions based on assumptions, biases, or incomplete information, which can lead to errors.
> 2. Fairness: In a situation involving others, understanding the facts first ensures that your judgment is fair and unbiased. This is especially important in conflict resolution, disciplinary action, or legal decisions.
> 3. Confidence: When your judgment is based on factual information, you can be more confident in your decision-making. This can help you to act decisively and persuasively.
> 4. Credibility: People tend to trust judgments that are clearly based on facts. Therefore, taking the time to understand the facts can enhance your credibility, whether you're a leader, a team member, or a friend.

5. Learning and Growth: The process of gathering and understanding facts can be a learning experience in itself. It can broaden your perspective, deepen your understanding, and enhance your problem-solving skills.
6. Risk Mitigation: Understanding the facts before making a judgment helps to identify potential risks and challenges, thereby allowing for better risk mitigation strategies.

At this stage the reader may be wondering how this is relevant to RCM!

The relevance is that Failure Effects (question 4) describe what happens. They are the facts.

Failure consequences (question 5) considers the extent to which the failure matters. This is a judgement!

3.4 The complete Information Worksheet

The first four questions of RCM are represented as four columns on a worksheet known as the Information Worksheet. This is a Functional Failure Modes and Effects Analysis (FMEA), because it has the important benefit of functions that classify the failures.

Function	Failure Mode	Failure Mechanism	Failure Effect
1 To received and separate distillate from its impurities at a rate of up to 3 tonnes per hour.	A Unable to receive and separate distillate from its impurities at a rate of up to 3 tonnes per hour.	1 Inlet blocked by foreign object	The filling side level remains static. Eventually the operator will become aware that the drum has not filled. Upstream flow indication is zero. Upstream process impacts can be minimised by re-routing. Downstream production or supply to boilers lost. Clear the blocked inlet - this may take between 4 hours and 4 days; personal protective equipment (PPE) needed.
1	A	1 Inlet line joint or gasket split (most likely following maintenance intervention	Distillate may escape to the oily sewer - effluent will probably go out of specification. Shift Technician may notice odour in vicinity of the leak. Upstream process impacts can be minimised by re-routing. Downstream production or supply to boilers lost. Possible fire hazard - if a fire results, then secondary damage may be caused to surrounding assets. Downtime to replace gasket 8 hours - PPE required.

Figure 19: Example Information (FMEA) Worksheet

The next stage of RCM is to decide what should be done to manage each failure through the application of the Decision Diagram which incorporates the remaining three questions.

CHAPTER 4

Failure Consequences

The 7 master questions of RCM are shown below. This chapter addresses the fifth of these – Failure Consequences

1. What do we want the asset to do? – Functions
2. How can it fail? – Functional Failures
3. What causes the Functional Failures? – Failure Modes
4. What happens when a failure occurs? – Failure Effects
5. How much does each failure matter? – Consequences
6. Can we predict or prevent failure and should we be doing so? – Proactive Tasks
7. How should we manage the failure if prediction or prevention is not an option? – Default Tasks

What is meant by failure consequences in RCM?

The failure Consequence is a judgement about the extent to which the failure matters. Failure Modes are classified into different categories of consequence by asking a series of questions.

Why do we consider failure consequences in RCM?

We are only bothered about failures because they matter in some way – because they have a consequence. If the failure does not matter much, we would not want to spend much time and effort managing it because it would not be worth it. Consideration of consequences is therefore a key aspect of deciding whether we manage the failure mode in any way.

This is important because use of scarce resources needs to be justified. If the failure is serious we can justify spending time and effort preventing, reducing or eliminating the consequences. If a failure does not matter, such investment may not be justified. In fact, unless a maintenance task changes the consequences of failure in some way, there will be no point doing it – for example, if the cost of leaving a bearing to fail is the same as the cost of preventing/ predicting the failure, there is no point preventing/ predicting the failure.

How do we characterise consequences?

The consequences of failure are determined by selecting one of the four columns in the Decision Diagram, by asking the questions at the top of the diagram and following the flow. This is done for each failure mode.

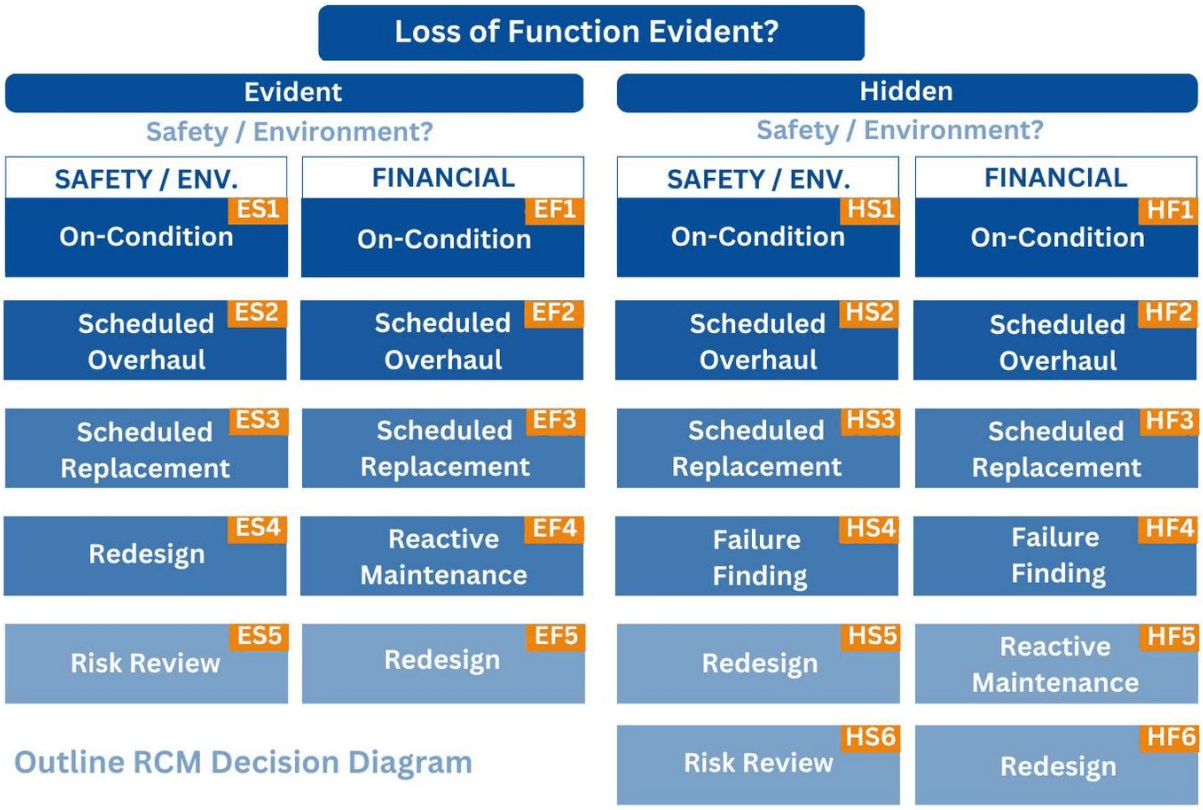

Figure 20: Outline RCM Decision Diagram

4.1 Evident and Hidden Failures

The first question separates failures into two broad classes of failure consequences – Evident and Hidden, by asking the question:

Will the loss of function or secondary damage caused by this single failure eventually become evident to the organization under normal system operating conditions?

- 'Yes' means failure is Evident
- 'No' means failure is Hidden

Why distinguish between hidden and evident failures?

This distinction is vital because:

- We will have different choices to make when it comes to deciding how to manage hidden failures.
- There are no immediate consequences for hidden failures - we need two or more things to happen before we would be affected, so managing hidden failure is more complicated

To illustrate these points, consider the failures on the left-hand column of the table below. If a smoke detector, airbag or PRV fail on their own (in certain ways) there will be no consequences. It is only if there is another independent failure – the failures in the middle column - that there are any consequences (these are given in the right-hand column).

Hidden Failure	Failure it protects againts	Multiple Failure
Smoke Detector failed	Smoke and Fire	Conflagration, with risk to human life
Airbag failed	Impact with steering wheel and fascia	Increased severity of injury or death
PRV failed	Overpressure	Explosion, injury or death

Figure 21: Hidden Failures only matter if there is another failure or adverse event

The devices on the left-hand side are also known as protective devices because they protect against another failure or mitigate against the consequences of failures. However, just because protective devices are fitted does not mean that the system can be ignored. Protective devices can also fail and if they are in a failed state when needed, a multiple failure would result. Multiple failures can sometimes be catastrophic, so careful management (including maintenance) of the different elements of the system may need to be considered.

4.1.1 Important points about the hidden/evident question

There are some important words/ terms in the hidden/ evident question that need to be expanded upon. These are explained below.

Single Failure

Each failure is considered on its own, assuming everything else is working, so for the PRV, Airbag and Smoke Detector above, the failure of these devices alone is considered (i.e. there is no fire, crash or overpressure (because these are separate failures)).

The Question of Time

A failure cannot be described as hidden because it takes a long time to become evident. If a failure will eventually become evident then it counts as evident. For example, if a water pipe corrodes over many years or decades leading to a leak which would cause loss of supply, this single failure would eventually become evident, even though it takes a long time to do so. (Note that the loss of supply etc are all effects of the initial, single failure).

Failures may become evident either via instrumentation or alarms and annunciation in a control room, by direct report from operational staff or by report from third parties (e.g. members of the public).

Loss of Function

The question is not asking whether people will know what the problem is, but whether they will know there is a problem. For example, if a family car breaks down because of a blocked fuel line, the driver will know that function has been lost, even though the driver may not know why. Therefore, the answer to the question would be 'yes' and the failure would be judged as evident.

Normal System Operating Conditions

When asking this question, it must be assumed that no maintenance is carried out, such as testing the device. The question is asking whether the loss of function will become evident to the organization without any maintenance intervention (the questions about maintenance come later in the Decision Diagram).

4.1.2 Some examples of hidden failures

There can be many hidden failures on industrial systems. Here are a few more examples:

1. Fire Suppression Systems: These systems are crucial for safety in facilities but are typically only activated when a fire is detected. If a component of the fire suppression system fails (e.g., a sprinkler head gets blocked), the failure might go unnoticed until the system is needed during a fire, at which point the failure could have disastrous consequences.
2. Backup Generators (or any backup or redundant equipment): A backup generator is a classic example of a system prone to hidden failures. It is not in constant use, so if it fails due to an issue like a faulty fuel pump or battery, the failure might not be noticed until there is a power outage and the generator doesn't start.
3. Redundant Systems: Redundant systems are designed to take over if the primary system fails. However, if the redundant system itself fails and this is not detected, it could lead to a situation where there is no backup when the primary system fails.
4. Safety Valves (PSVs): In pressure systems, safety valves are designed to release pressure if it exceeds a certain limit to prevent a catastrophic failure. If the valve is stuck or otherwise fails, this might not be apparent until an overpressure situation occurs and the valve fails to operate.
5. Alarm (and trip) Systems: Alarm systems are designed to warn operators of abnormal conditions. However, if the alarm system fails (for instance, a sensor fails), this will usually not be noticed until an abnormal condition occurs and the alarm fails to sound.

Other terms for hidden failures

Hidden failures are also known as unrevealed or dormant failures.

4.1.3 Further points about hidden failures

Different failure modes of protective devices

Many protective devices fail in two ways:

1. By failing to operate when needed
2. By operating when they are not needed

The first of these is a hidden failure. For example, if a sprinkler system does not operate when there is a fire, this is a hidden failure.

The second of these is an evident failure – if the sprinkler system goes off when there is no fire, this would become evident to the organisation!

In an RCM analysis, this would be handled as two different failure modes for the low-level alarm (at 120 000 litres) in the diagram below.

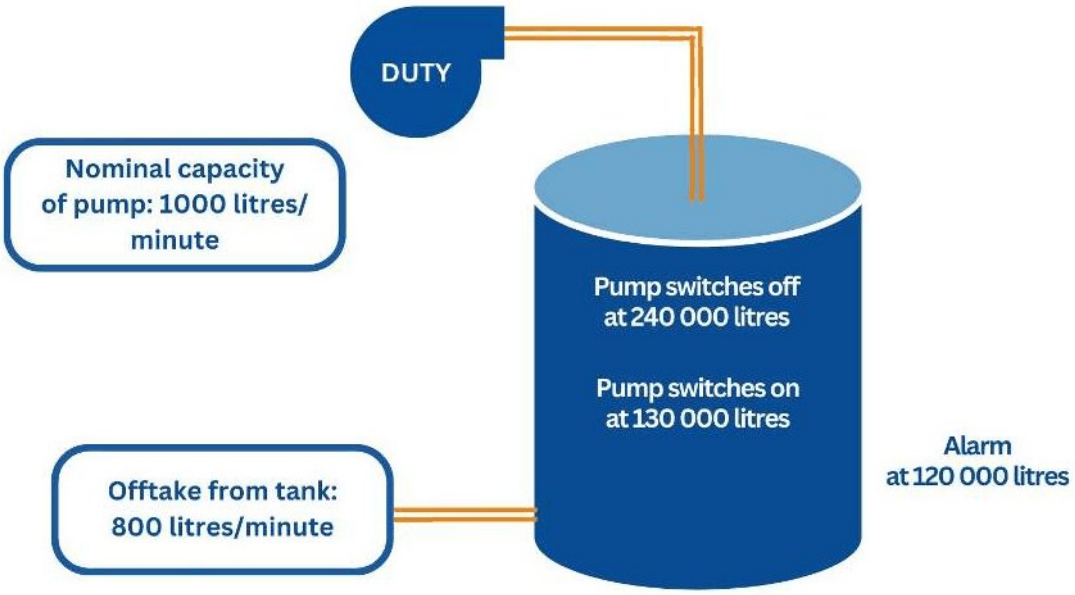

Function	Functional Failure	Failure Mode	Failure Effect
To alarm if the water level falls to 120 000 litres	Fails to alarm if the water level falls to 120 000 litres	Low level alarm fails closed	This only matters if the water level falls to 120 000 litres because of some other failure such as the pump, or the control switch. The water level would fall below 120 000 litres and there would be no alarm. The tank would run dry and the process would stop. Time to repair 4 hours.
	Alarms if the water level does not fall to 120 000 litres	Low level alarm fails open	There would be an alarm in the control room during normal operation. Operations would need to investigate the cause of the alarm. A local gauge could be used to determine that the water level is within the desired range, and this would point to a spurious operation of the alarm. Time to repair 4 hours.

Fail safe devices

In RCM, a fail-safe device is a protective device that becomes evident to the organisation if it fails on its own. In the context of the discussion about hidden failures, it is a protective device that fails in an evident manner. This means that the failure of the device will then be economic, but it will not affect safety (hence the term 'fail-safe').

Some typical examples of fail-safe devices are given below:

1. Industrial Fire or gas detection systems: For many failure modes of such systems, a fault signal may be sent to a central monitoring station or control panel, making the failure evident.
2. Industrial Safety Systems: Many industrial safety systems, such as machinery safety interlocks, are designed to alert operators or shut down equipment if a dangerous condition is detected. If these systems fail, they may also provide an alert or shut down equipment as a precaution, making the failure evident.
3. Uninterruptible Power Supply (UPS) Systems: These devices provide emergency power to a system when the input power source or mains power fails. Many UPS systems have alarms or indicators to alert users when the UPS itself has failed.
4. Intrusion Detection Systems: These systems are designed to detect unauthorized entry into a building or area. If the system fails, it can often send a fault signal to a central monitoring station, making the failure evident.

5. Health Monitoring Systems in Aircraft: Aircraft often have systems for monitoring the health of various components. If a component such as an engine monitoring system or flight control computer fails, it can trigger a cockpit warning or even an automatic system response, making the failure evident.
6. Safety Instrumented Systems (SIS) in Process Industries: These systems are designed to bring a process to a safe state in case of abnormal conditions. They usually have self-diagnostics and will alert if any part of the SIS fails.

In these examples, the protective devices are designed to alert the organization in the event of their own failure.

Self-checking systems

Most of the devices listed above have systems that check themselves to make hidden

failures evident. Another example is modern smoke and flame detectors which often come equipped with self-checking or self-test systems that enable them to ensure they are functioning properly and alert you if there is an issue. Here is an overview of these features:

1. Battery Check: Smoke detectors and flame detectors will often have a low battery warning, typically an intermittent chirp or beep, to alert you when the battery level is low and needs to be replaced.
2. Sensitivity Check: Some detectors also come equipped with a self-checking feature to monitor their own sensitivity. Over time, smoke detectors can accumulate dust and debris, which can affect their sensitivity to smoke. A self-check feature will alert you if the sensitivity has dropped below a certain level, indicating it may be time for cleaning or replacement.
3. Power Supply Check: If the detectors are hardwired to the electrical system, they may have a feature to check for proper power supply. If the power supply is interrupted, a warning will be issued.
4. Sensor Health Check: More advanced models may even be able to monitor the health and function of the smoke or flame detection sensors themselves. If the sensors fail or degrade over time, the detector can alert you that it's time for a replacement.
5. Self-Test Button: Most domestic smoke detectors will have a self-test button that you can press to ensure the alarm is working correctly. When pressed, the detector should emit a loud alarm sound.
6. Interconnectivity Test: For systems where multiple detectors are connected, a self-checking feature might test whether the interconnection is functioning correctly, i.e., when one detector activates due to smoke or flame detection, it should trigger all other connected detectors.

These self-checking features are crucial in ensuring the detector is in working order and can correctly warn occupants in the event of a fire.

A caution about self-checking and fail-safe devices

In general, self-checking devices are good news. Hidden failures increase risk of sometimes catastrophic multiple failures - but we do not know they have occurred. Self-checking devices make such failures evident, effectively by testing automatically. However, there are two cautions about such devices:

1. The self-checking functionality can fail
2. The self-checking coverage can be incomplete, and there can still be hidden failures lurking in the system

To give a simple example to illustrate the second point – some domestic smoke detectors have:

- a red test button to test the electronics and the siren
- a check for low battery voltage which sounds the alarm when the voltage is low (usually in the middle of the night!)

However, neither of these arrangements will test the sensor, which can fail in a hidden manner. In many protective systems the sensor is often the least reliable component because it is the part of the system most exposed to the outside environment, but it is the part of the system that is often the most difficult to fully test. To give some examples from the aircraft industry, the following crashes involved sensor failure (note that these were not necessarily protective devices):

1. Air France Flight 447 (2009): This flight from Rio de Janeiro to Paris crashed into the Atlantic Ocean, killing all 228 people on board. The accident was largely attributed to the failure of the aircraft's pitot tubes, which are sensors used to measure airspeed. These tubes became blocked by ice crystals, causing inconsistent airspeed readings and confusion among the pilots. This event triggered a series of errors in the automated and manual flight controls, leading to the aircraft stalling and eventually crashing into the ocean.
2. Lion Air Flight 610 (2018) and Ethiopian Airlines Flight 302 (2019): These two Boeing 737 Max crashes, which collectively resulted in 346 fatalities, were attributed in part to a faulty angle of attack (AoA) sensor. In both incidents, the erroneous data from the AoA sensor activated the Maneuvering Characteristics Augmentation System (MCAS), an automatic flight control system. The MCAS repeatedly pushed the nose of the plane down, leading to a fatal dive. These two crashes led to the worldwide grounding of the Boeing 737 Max.
3. Aeroperú Flight 603 (1996): This flight crashed into the Pacific Ocean, resulting in the death of all 70 people on board. The accident was caused by a maintenance error in which protective tape was accidentally left over the static ports (used to measure altitude and airspeed) after the aircraft was cleaned. The covered ports provided erroneous readings to the cockpit instruments, leaving the crew confused and unable to control the flight properly.

Please note that while sensor failures contributed to these crashes, they are usually just one factor in a complex sequence of events. In many cases, these incidents also revealed problems with aircraft design, operator training, maintenance procedures, or regulatory oversight, which were addressed in the subsequent investigations and safety recommendations.

Hidden failures and modern Industrial Instrumentation

Overall, the trend in modern industrial instrumentation has been to replace traditional mechanical switches have with analogue and digital devices. Here are a few examples:

1. Pressure Switches: In some applications, mechanical pressure switches have been replaced by pressure transmitters that provide an analog output. This allows for more precise control and monitoring of system pressures.
2. Level Switches: In tanks or silos, old-style float switches used for level detection have been replaced by modern devices like ultrasonic or radar level sensors, which can provide continuous level measurement and more accurate control.
3. Flow Control: Mechanical flow switches, used to detect the presence or absence of flow in a system, are increasingly being replaced by digital flow meters or transmitters. These devices can provide continuous measurement of flow rates and are often more accurate and reliable.

4. Proximity Switches: Old-style mechanical limit switches used for position detection are being replaced by inductive or capacitive proximity sensors in many applications. These can detect the presence of an object without physical contact, reducing wear and tear.
5. Safety Systems: Mechanical safety switches used for emergency stop functions are being supplemented with light curtains or laser scanners that can provide a safer and more flexible protective zone around hazardous machinery.

In such systems, a higher proportion of failure modes may become evident because the devices may not 'stick' in the traditional way, and total failure of the device or the control system would usually become evident.

RCM and Modern Instrumented Systems Although it was developed over 50 years ago the RCM process is more relevant than ever for modern equipment incorporating smart sensors for the following reasons:

1. Increased Complexity: Modern systems are more complex than ever, with intricate interactions between various subsystems. However, these interactions need to be understood between users and maintainers. RCM, with its structured and systematic approach, helps manage this complexity. RCM is also a very powerful learning tool and it is essential for safe and reliable operation that the people operating and maintaining equipment understand how it works. This means that an RCM analysis of a complex system may get challenging, but that is the fault of the system designers, not of the RCM process!
2. Data-driven Decision Making: Smart sensors provide a wealth of data, which can be harnessed to monitor equipment performance and predict failures. This aligns well with RCM's focus on predictive maintenance based on actual equipment condition. In many cases, organisations collect a lot of data about the equipment but they don't use it. RCM provides a means to use the data collected to make a difference.
3. Integration of Technologies: The convergence of technologies like the Internet of Things (IoT), Artificial Intelligence (AI), and machine learning in the industrial space allows for enhanced predictive capabilities. This advances the application of RCM, providing tools to better anticipate and mitigate failures.
4. Safety and Compliance: Modern equipment and operations often face stringent safety regulations. RCM supports this by emphasizing the safety and operational consequences of equipment failures and providing the audit train for their proper management.
5. Human and external factors: Clever and sophisticated systems can still fail because of 'simple' things. For example, the external environment can cause failures. In addition, the interaction of humans with such systems can cause failures. These factors are considered in RCM analyses of such systems.

4.2 Safety and Environmental Consequences

If the failure is evident the next question that is asked is:

Will the loss of function or secondary damage caused by this Functional Failure have an intolerable adverse effect on Operating Safety or the Environment?

An important phrase in this question is the words 'intolerable adverse effect'. This is to deal with the issue of whether we should include every failure with the possibility of a hazard under Safety Consequences. For example, a band joint on 11kV switchgear suffers ingress of moisture and disrupts – this is potentially dangerous but very very rarely results in injury, so this may not be judged a safety consequence. To be judged a safety consequence the

safety risk must be above a tolerable threshold. The key point here is that the tolerable risk is not zero! The section below explains what the term 'tolerable risk' means.

4.2.1 Tolerable Risk

Most things we do carry a measure of risk which we tolerate knowingly or otherwise. A good reference on this subject is the UK Health and Safety Executive philosophy document from 2001 entitled "Reducing Risks, Protecting People". (Ref) The document states that "In this context tolerable does not mean acceptable. It refers instead to a willingness to live with a risk so as to secure certain benefits and in the confidence that the risk is one that is worth taking and that it is being properly controlled."

Virtually every activity we carry out involves a measure of risk. These activities are carried out willingly and the risk is tolerated. The level of risk associated with these activities can be used to assist with making risk-based decisions at work.

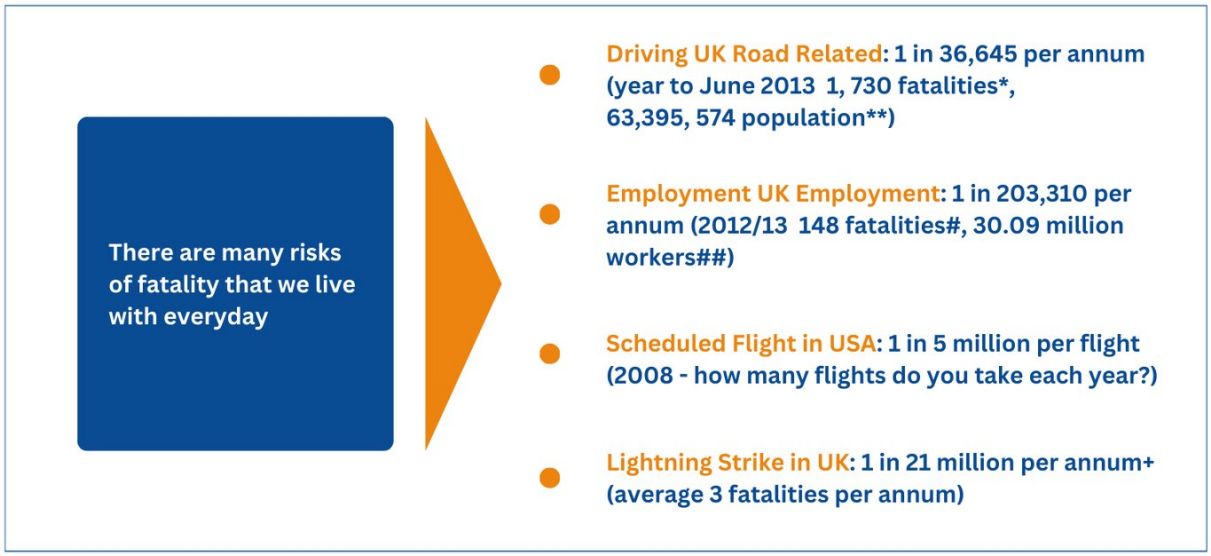

The table below shows some risks that are tolerated in some everyday activities – the figures are the risk of a fatality per person per year.

Figure 22: Everyday risks of fatality

The table below shows the risks associated with driving (data from 2002) – one of the riskiest activities we carry out (but a risk we (invariably) tolerate).

Road deaths	per 100,000 popn.	per 100,000 vehicles
UK	7.6	15
Germany (W)	11.3	19
Switzerland	12.2	21
Turkey	13.9	21
Italy	13.9	24
Belgium	16.7	35
France	17.3	36
Austria	17.8	36
Spain	20.2	47
Germany (E)	21.1	48
Portugal	34.4	86

Figure 23: 2002 European statistics for road deaths per 100,000 population / vehicles

While the table below shows some figures for certain activities in the USA (Sources: National Safety Council, National Academies, The Economist). (Ref)

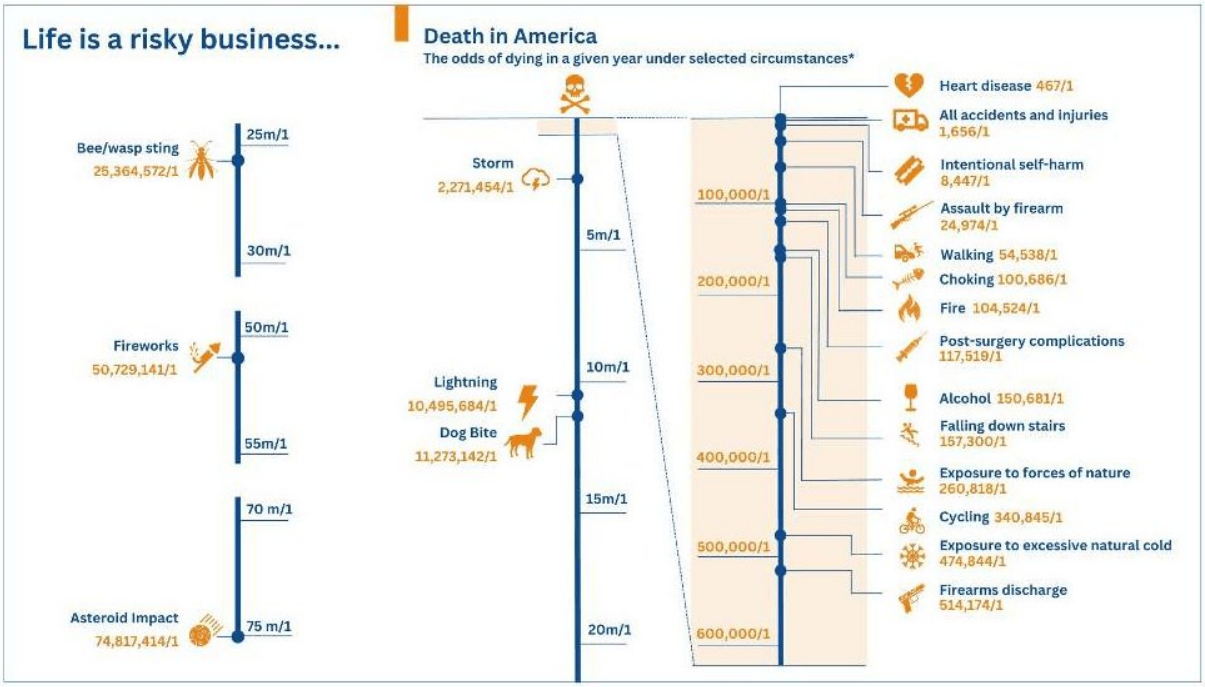

Figure 24: USA Comparative fatality risks

The value of these is that it is possible to use figures such as these to provide a framework to make sensible risk based decisions in maintenance and design (as will be shown later in this document).

4.2.2 Why decides about tolerable risk and what influences their decision?

Decisions about risk should be taken by those directly facing the risk, rather than a third party. However, the perceptions of customers, neighbours and society as a whole must also

be addressed and an organization has a moral and legal responsibility to assess and work within these parameters.

Decisions about what levels of risk are tolerable are not always objective and consistent. For instance, statistically, flying is much safer than driving yet some people would rather drive than fly. One factor that influences people's perception of risk is the degree of control they feel they have – when people fly they have less control. Another factor is the degree of choice – people often tolerate higher levels of risk in situations of their choice (diving, climbing, kite surfing etc) - than they would be prepared to accept if they had to do such activities as part of their jobs.

4.2.3 Safety consequences – who decides?

The people who are exposed to the risks and who suffer the consequences would be involved in deciding whether a failure mode is classified as a safety consequence. It would be unethical for outsiders to make safety decisions without involving the people who are exposed to those risks, or at the very least, without their informed consent. This is another reason why RCM analyses are carried out by teams of people who know the asset best.

4.2.4 Environmental Consequences

Society's view of environmental damage has become increasingly less tolerant in recent decades. Environmental issues and incidents take many forms and include Oil Tanker Disasters, the Disposal of Oil Platforms and Agricultural Pollution (phosphates in water and pesticides).

The legal power and focus of government agencies is increasing, there is more frequent Monitoring / Reporting (there is an onus on organisational openness) and fines are more significant. In addition, there is potential for Public Relations damage and share price effects of environmental excursions.

Furthermore, Environmental consequences may eventually become human safety issues. Famous examples include:

- Minimata (mercury waste converted to dimethyl mercury in shellfish, Japan 1960s)
- Pacific Gas & Electric (Use of Cr6 as anti-corrosion agent at gas pumping stations
  - which was portrayed in the film Erin Brokovich)
- Love Canal (housing built on toxic waste dump in NY State, 1970s)

In extreme cases, companies that breach environmental standards may be forced to shut down.

Consequently, Environmental consequences are considered in the same way as safety consequences in the RCM process.

The following might be included under Environmental Consequences:

- Oil or Chemical spillage
- Out of Consent discharge (Turbidity, pH)
- Leakage from fluid-filled cables
- Noxious or poisonous discharge

Environmental Standard or regulation

An important point about the Environment question in RCM is that a failure is classed as Environmental if it breaches 'any known standard or regulation.'

4.2.5 Is it Safety / Environment?

If the risk to safety/ environmental would exceed tolerable levels:

- Then the answer is "Yes" – we consider it as a Safety / Environmental consequence and we need to manage the probability of failure

If the risk is within tolerable levels:

- Then the answer is "No" – we consider it as a Financial consequence and how to manage the costs associated with failure

4.3 Financial Consequences

Financial consequences include loss of, or reduced production/supply, quality or performance problems, loss of contract / compensation, higher operational costs, repair costs (resources, spares), secondary damage or fines or other penalties.

Some failures only have repair and secondary damage to consider such as loss of a duty pump where there is a standby pump that cuts in automatically, or open circuit failure of one out of four Primary 33 – 11kV transformers, where there is "firm" capacity (peak demand can still be met in the event of the largest transformer being unavailable). Note that does not mean that such failures will always default to a policy or run-to-failure – in some cases the costs of secondary damage are significant and maintenance tasks are cost effective to prevent these costs. In other cases, for critical systems with high consequences, even if an item is protected by a redundant device it may still be cost effective to carry out maintenance on the duty device to improve its reliability, and therefore the reliability of the system as a whole.

Sometimes the point is made that safety/ environmental consequences cost money and that financial factors should be considered during the evaluation of such failures. In fact this is not against the principles of RCM – it is good practice to consider the economics of failure with safety/ environmental consequences – it is just that the probability of failure is the main driver of the decision, because the safety/ environmental consequence is more important. Provided a failure management decision reduces the probability of failure to a tolerable level, it is of course desirable to select a solution that is the most economic. In many cases, managing a risk to reduce its probability to a tolerable level also of course has the benefit of improving the economics (since safety/ environmental consequences are often also extremely expensive).

4.4 Evaluating safety consequences of Hidden Failures

Like Evident Failures, Hidden Failures belong to one of two Consequence categories:

- Safety/Environment
- Financial

If the failure is hidden the question we ask is:

Will the loss of function or secondary damage associated with the multiple failure have an intolerable adverse effect on Operating Safety or the Environment?

When we ask the question we need to consider whether the risk of the multiple failure exceeds tolerable levels.

If the risk exceeds tolerable levels then the answer is "Yes" – we consider it as a Safety / Environmental consequence and we need to manage the probability of the multiple failure. If we have a multiple failure - a hidden failure and the failure it protects against - the ultimate consequences could be major and could include incidents such as an explosion because of overpressure and a failed PSV or a more serious fire because a fire alarm or smoke detector does not work.

If the risk is within tolerable levels then the answer is "No" – we consider it as a Financial consequence and how to manage the costs associated with multiple failure – examples include cost of lost production if standby pump fails to operate when duty pump fails or the cost of a new electric motor when overload trip fails and is needed. For the examples of multiple failures listed below, the first three would probably be judged as hidden safety, the second two hidden economic.

Hidden Failure	Failure it protects againts	Multiple Failure
Smoke Detector failed	Smoke and Fire	Conflagration, with risk to human life
Airbag failed	Impact with steering wheel and fascia	Increased severity of injury or death
PRV Failed	Overpressure	Explosion, injury or death
Spell checker failed	Spelling mistake	Spelling mistake in email
Low level alarm switch failed	Pump failure	Increased in lost production because of no alarm

CHAPTER 5

Proactive Maintenance Tasks

The 7 master questions of RCM are represented below. This chapter addresses the sixth of these – can we predict of prevent failure and should we be doing so?

1. What do we want the asset to do? – Functions
2. How can it fail? – Functional Failures
3. What causes the Functional Failures? – Failure Modes
4. What happens when a failure occurs? – Failure Effects
5. How much does each failure matter? – Consequences
6. Can we predict or prevent failure and should we be doing so? – Proactive Tasks
7. How should we manage the failure if prediction or prevention is not an option? – Default Tasks

Having decided how much each Functional Failure matters, we can now decide how best to manage them. This is where we select the most appropriate failure management policy.

In this chapter will now consider what happens down each of the columns of the decision diagram as an appropriate failure management policy is sought.

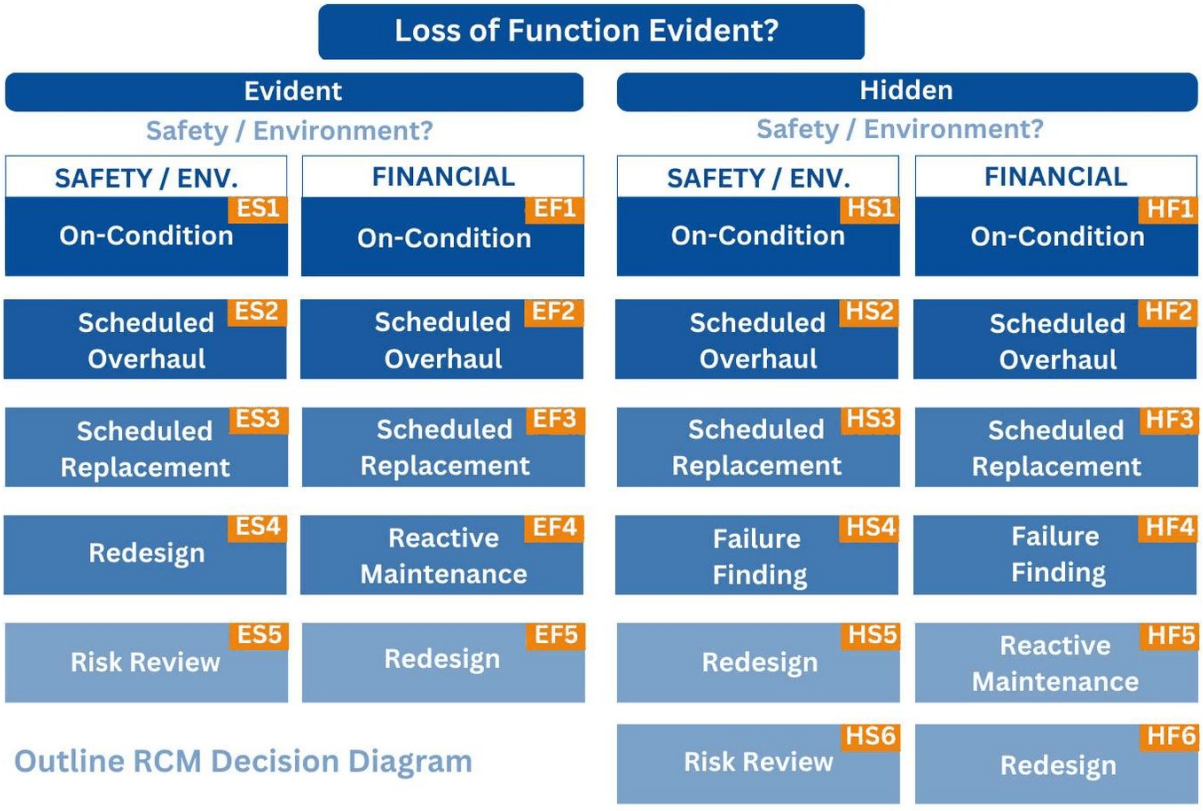

Figure 25: Outline RCM Decision Diagram

5.1 Technically Feasible and Worth Doing

Whatever the Consequence of the failure, the proposed method of managing that failure must be both:

- Technically feasible
- Worth doing

If a task is not technically feasible it will not be valid, but a task would not be appropriate if it was technically feasible but not worth doing. The RCM decision diagram asks questions about both criteria for each failure management task box – for example:

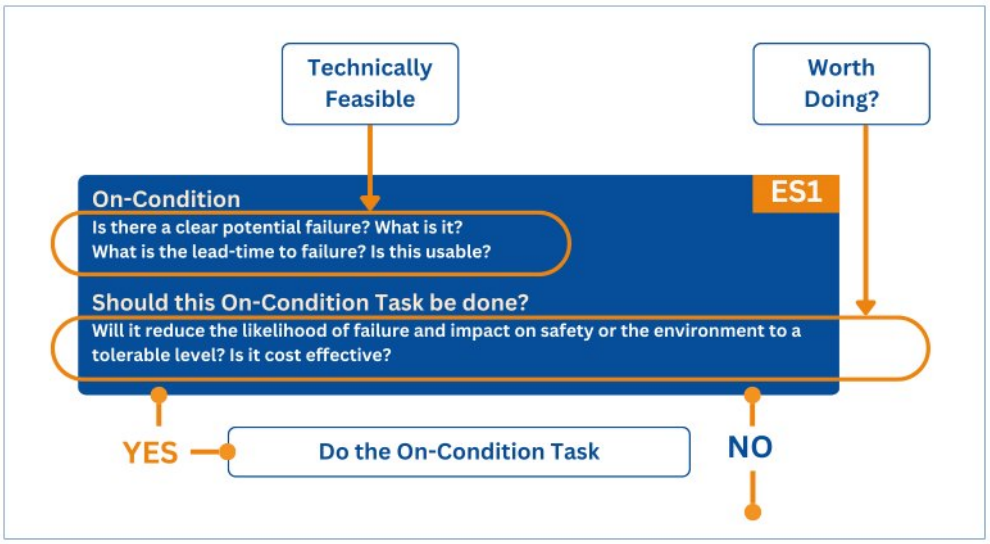

Figure 26: On-condition task evaluation – technically feasible and worth doing?

The following sub-sections consider the "technically feasible?" and "worth doing?" questions in more detail.

5.1.1 Technically Feasible

A failure management solution is technically feasible if it manages the failure to an acceptable degree. This means that must be possible to do the task, the task must minimize the likelihood of the failure and there should not be a significant risk it could fail before the next maintenance.

5.1.2 On-condition Maintenance

This is also known as condition-based or predictive maintenance and is the first type of maintenance task the decision diagram considers.

Condition-based maintenance is based on the fact that we often get warning signs that assets are deteriorating toward loss of functional and an on-condition task may be carried out to predict that a failure is occurring or about to occur. This will allow action to be taken to avoid or minimize the consequences of the failure. These warnings are also known as potential failures. If nothing is done then the failed state – a functional failure - will result. The usefulness of these warning signs depends to a large extent on how close they are to failure.

The diagram below shows a generic representation of progression towards failure.

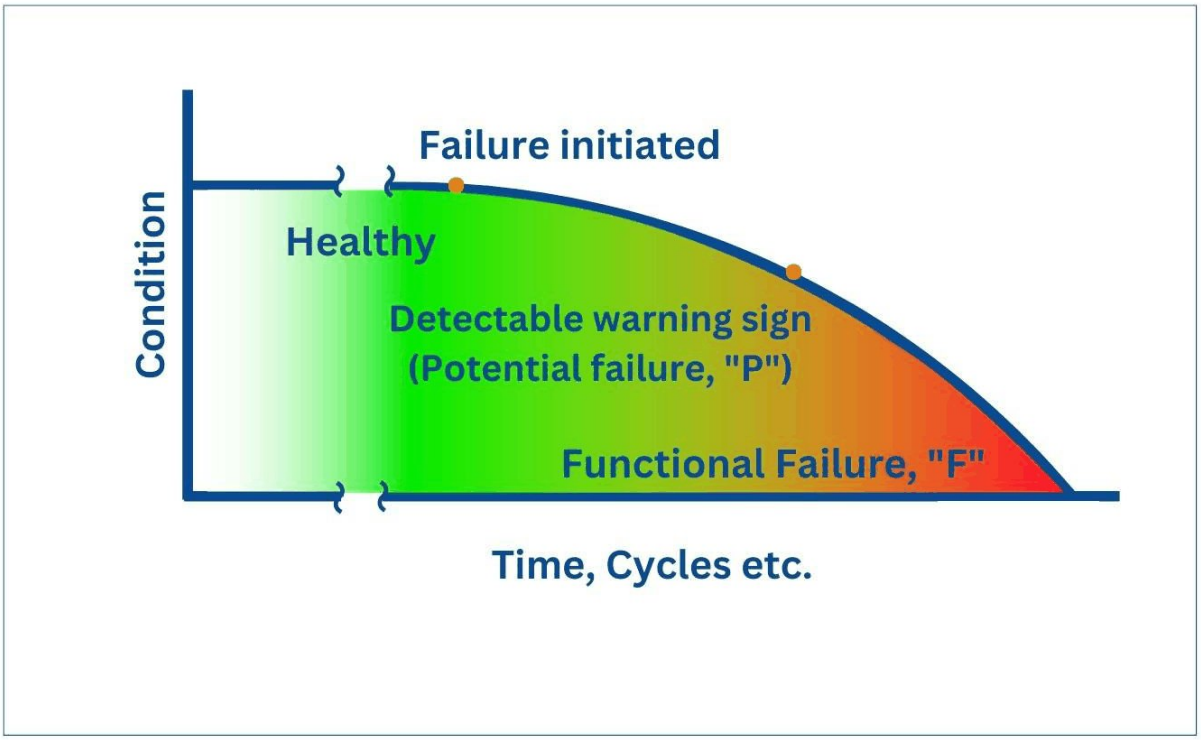

Figure 27: The “P-F” Curve

The frequency of a condition-based task is dictated by the time from a potential failure to a functional failure, also known as the p-f interval, or the lead-time to failure. To be technically feasible an on-condition task must identify a potential failure that gives clear warning that failure is in the process of occurring, gives enough time to react to avoid the consequences of failure and is usable (suitable for the organization).

To ensure that we don’t miss the warning sign we need to carry out the maintenance at an interval less than the lead time to failure. If the failure has Safety/ Environmental consequences the task is usually carried out at half the lead time to failure or less. If the failure has Financial consequences the task interval needs to be something less than the lead time to failure.

Examples of potential failures include cracks in structures, wall thinning in pipes, reduced performance of a gas turbine, compressor or pump and noise, heat, vibration and oil debris in bearings. In fact, for many failures there are often numerous potential failure conditions, and therefore numerous possibilities for condition-based maintenance tasks.

Note that the existence of the Condition Deterioration Curve may be entirely independent of age. In other words, once started, the rate of deterioration is similar for a new asset or one that is one year old or ten years old.

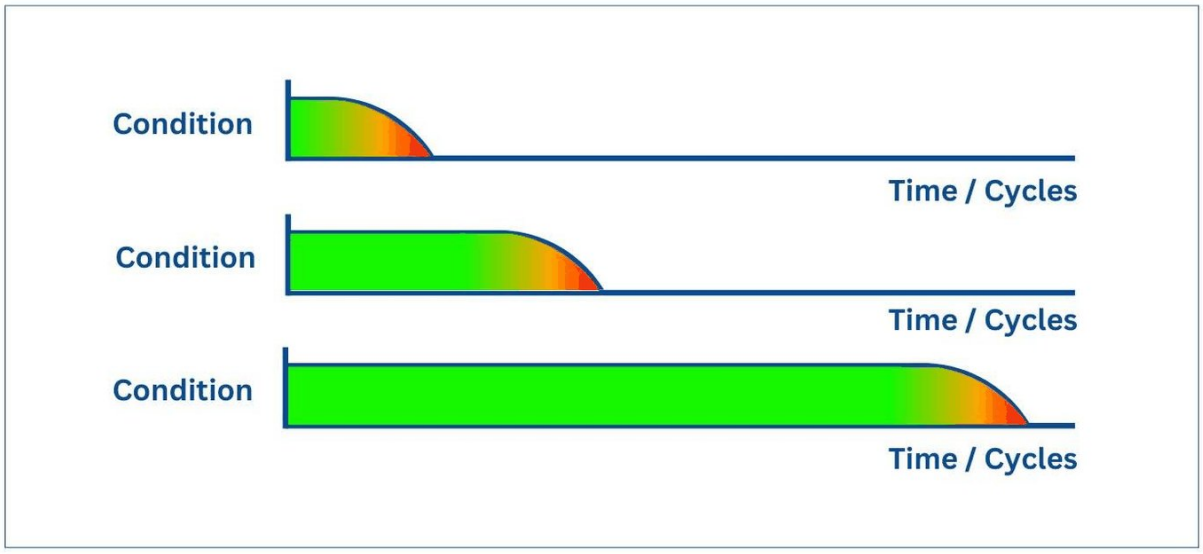

Figure 28: P-F may not be related to age at all

The questions for technical feasibility of condition-based maintenance ask:

- Is there a clear Potential failure condition?
- What is the lead-time to failure?
- Is the lead-time usable (i.e. can we inspect within the lead-time)?

5.1.2.1 Categories of Condition based task

There is a massive range of on-condition tasks which reflect the numerous potential failure conditions that can occur. Categories of on-condition tasks are given in the table below.

Condition-based Maintenance	
Category	Definition
Condition Monitoring	The use of specialized equipment to monitor the condition of other equipment.
Product Quality Monitoring	Monitoring the condition of the product to predict problems with the machine that makes the product
Performance Monitoring	Monitoring the performance of a machine to determine potential failures using external or build in gauges or other measurement devices
Human Senses	The use of the human senses (look, listen, smell, touch) to identify potential failure conditions

Condition Monitoring - the first of these categories – can be further sub-divided into six categories:

Condition Monitoring Category Definition

Dynamic Effects Monitoring	Looking for potential failures in the form of waves and pulses. This includes vibration monitoring.
Particle Effects Monitoring	Looking for potential failures in the form of particles (usually in oil)

Chemical Effects Monitoring	Looking for potential failures in the form of chemical changes to oil or chemicals in oil, the air or products
Physical Effects Monitoring	Looking for potential failures in the form of changes to the physical characteristics of equipment in the form of dimensional changes, cracks and viscosity changes
Temperature Effects Monitoring	Looking for potential failures in the form of temperature changes or thermal effects
Electrical Effects Monitoring	Looking for potential failures in the form of changes to the electrical characteristics of equipment

5.1.2.2 Examples of some common condition-monitoring and condition-based maintenance techniques

The table below gives more detail about some of the most widely-used Condition Monitoring Techniques, including common applications, required skills, advantages and disadvantages and typical P-F intervals.

Note that for virtually all the techniques, required skills include knowledge of the equipment being assessed, how to carry out the test properly and how to interpret the results. For example, dirty oil in one application may not be a problem in another.

CBM Strategy	Common Applications and principle	Required Skills	Advantages	Disadvantages	Typical PF-Interval
Vibration Analysis	Detecting defects in rotating equipment through changes in vibration characteristics	Knowledge of vibration analysis equipment and data interpretation	Early detection of faults, minimizes downtime, improves equipment reliability	Costly equipment and expertise required, can be complex although modern software simplifies interpretation to some extent	Weeks to months
Oil Analysis	Monitoring: 1. lubricant properties or 2. External contaminants in oil or 3. Wear particles in oil	Knowledge of lubricants and oil analysis techniques	Early detection of equipment failures, improved equipment reliability, and longevity	Costly equipment and expertise required, limited diagnostic information, time-consuming sample collection and analysis	Weeks to months
Infrared Thermography	Detecting thermal anomalies in electrical, mechanical and static equipment using specialised thermal cameras	Knowledge of thermography equipment and data interpretation	Early detection of overheating components, improved safety, and equipment reliability	Limited diagnostic information, temperature readings affected by environmental conditions. Equipment can be expensive	Days to months

Ultrasonic Noise Testing	Detecting changes in sound waves to identify equipment defects. This may occur in pressure or vacuum systems and also electrical systems	Knowledge of ultrasonic equipment and data interpretation	Early detection of equipment failures, improved equipment reliability, and longevity	Limited diagnostic information, complex data interpretation	Weeks to months
Ultrasonic defect and thickness Testing	Measuring sub-surface defects or material thickness using high frequency sound	Knowledge of ultrasonic equipment and data interpretation	Early detection of equipment failures, improved equipment reliability, and longevity	Limited diagnostic information, complex data interpretation	Weeks to months
Radiography	Measuring sub-surface defects or material thickness using X-ray of gamma radiation	Knowledge of equipment and data interpretation	Early detection of equipment failures, improved equipment reliability, and longevity	Safety requirements need to be observed when using the equipment	Weeks to months
Cameras, boroscopes and endoscopes	Visual inspection internally, or of inaccessible equipment	Knowledge of equipment and data interpretation	Early detection of equipment failures, improved equipment reliability, and longevity	Skill required to use the equipment	Days to months

Chemical analysis techniques	Looking for chemical changes and detecting the presence of chemicals (including wear particles and contaminants (which includes water)). A range of techniques apply here including spectroscopy	Knowledge of equipment and data interpretation	Early detection of equipment failures, improved equipment reliability, and longevity	Expensive equipment for some techniques – Labs are often used.	Days to months
Motor Circuit Analysis	Detecting electrical faults in motors	Knowledge of electrical equipment and data interpretation	Early detection of motor faults, improved motor reliability	Costly equipment and expertise required, limited diagnostic information, time-consuming	Months
Acoustic Emission Testing	Detecting equipment defects through sound emissions	Knowledge of acoustic emission equipment and data interpretation	Early detection of equipment failures, minimizes downtime, improves equipment reliability	Limited diagnostic information, complex data interpretation	Months

Corrosion Monitoring	Detecting corrosion in metal structures. This can be through use of coupons or probes measuring electrical resistance, hydrogen, sand or biological parameters	Knowledge of corrosion monitoring equipment and data interpretation	Early detection of corrosion, minimizes equipment failures and downtime	Limited diagnostic information, costly equipment and expertise required	Months to years
Magnetic Flux Leakage Testing	Detecting metal defects and corrosion in pipelines and tanks	Knowledge of magnetic flux leakage equipment and data interpretation	Early detection of equipment failures, minimizes equipment downtime	Costly equipment and expertise required, complex data interpretation	Months
Eddy Current Testing	Detecting defects in conductive materials such as pipes and tubes	Knowledge of eddy current equipment and data interpretation	Early detection of equipment failures, improved equipment reliability and longevity	Costly equipment and expertise required, complex data interpretation	Months
Motor Current Signature Analysis	Detecting electrical faults in motors through current measurements	Knowledge of electrical equipment and data interpretation	Early detection of motor faults, improved motor reliability	Costly equipment and expertise required, limited diagnostic information	Months
Insulation resistance test	Testing insulation resistance on windings of motors, generators etc	Knowledge of equipment tested	Early detection of equipment failures, improved equipment reliability	Equipment needs to be taken offline to carry out the test	Months to years
Partial discharge testing	Insulation breakdown in a wide range of electrical equipment including switchgear, transformers, generators, cables	Knowledge of electrical equipment and data interpretation	Early detection of equipment failures, improved equipment reliability	Several data points needed. Cannot pinpoint location of problem. Standards difficult to define	Weeks to months

The table below provides similar details about other types of condition-based techniques (as distinct from Condition Monitoring)

CBM Strategy	Common Applications and principle	Required Skills	Advantages	Disadvantages	Typical PF-Interval
Visual Inspection	Inspecting equipment visually for defects or damage	Knowledge of equipment	Early detection of equipment failures, minimizes downtime, improves equipment reliability. Widely used. Low marginal cost if people already in attendance of the assets are used	Limited diagnostic information. May not detect internal defects. Subjective. Access may not be possible in some environments	Daily to weeks
Other human senses	Listening or feeling to detect equipment deterioration	Knowledge of equipment	Early detection of equipment failures, minimizes downtime, improves equipment reliability. Low marginal cost if people already in attendance of the assets are used	Limited diagnostic information. Subjective. Access may not be possible in some environments	Daily to weeks

Product quality monitoring	Identifying equipment deterioration by measuring the quality of the product made by the equipment	Knowledge of process and data interpretation	Early detection of wide range of equipment problems. Potentially a large number of applications. May use built-in sensors or sampling/ testing that is carried out anyway (although not necessarily for maintenance purposes)	Expertise required, limited diagnostic information	Days to Months
Performance Monitoring	Identifying equipment deterioration by measuring the performance of the asset	Knowledge of equipment and data interpretation	Potentially a large number of applications. May use built-in sensors.	Expertise required, limited diagnostic information	Days to Months

5.1.2.3 Condition-based Maintenance using calculated measures from one or more inputs

One notable area of condition-based maintenance is the principle of using a calculation from one or more inputs to determine if a failure is occurring or about to occur. For example:

- A heat exchanger efficiency loss can be calculated by considering inlet and outlet temperatures, fluid flow, fluid characteristics and ambient temperature, and comparing to a calculated value when the heat exchanger was new/ clean.
- Pump deterioration can be identified by taking relevant inputs to determine the flow in relation to the 'as-new' condition on the pump curve
- Deterioration of a valve can be identified by how long it takes to settle down in the event of a transient, or the number of overshoots it exhibits
- Deterioration of a duty pump protected by a standby can be identified by how often the stand-by cuts in There are many more examples that could be given and the number of opportunities in this area is almost limitless. Most of these examples would classify as performance monitoring, but some might be versions of condition-monitoring. In all cases, if done properly, the calculated measure can provide a more accurate and more robust measure of equipment condition than a simple, single measure. In order to carry out such calculations, more inputs are needed, but the trend in modern plants is to

provide more and more data. Indeed, many organisations collect lots of data, but don't use it to make decisions.

In some environments, potential failure conditions are referred to as 'indicators' and those involving calculations as 'calculated indicators'.

5.1.2.4 Manual Indicators

Manual indicators include traditional inspection methods that rely on human operators or maintainers to detect signs of wear or damage. Examples include visual inspection, sound and vibration analysis, and fluid analysis. For instance, technicians may visually inspect a bridge for cracks or corrosion, or mechanics may listen for unusual noises in a vehicle's engine. Their strengths include ease of use and low technology requirements. Their weaknesses are the reliance on human expertise and subjectivity.

5.1.2.5 Online Indicators

Online indicators, such as sensors and monitoring systems, provide real-time data on asset performance or condition. Examples include temperature, pressure, position, flow and vibration sensors. They would traditionally be built into a rule-based system that triggers an alert if the sensor data exceeds a specific threshold. These sensors enable real-time monitoring, continuous data collection, and quick anomaly detection. One important consideration in such systems is the alarm setting (when something is normally considered to be out of spec (ie failed)), and the pre-alarm setting (when something is failing).

5.1.2.6 Machine Learning-Based Indicators

Machine learning-based indicators utilize algorithms to analyze data and identify patterns and anomalies and predict equipment failure by analyzing historical data from various sensors. Unlike rule-based indicators, machine learning-based indicators can adapt to new situations and evolving data. This is a new area of technology. The beauty of RCM - a process that is over 50 years old - is that it is even more relevant with new maintenance techniques and developments. Machine learning algorithms need to preserve functions, they need to address failure modes, they need to be technically feasible and they need to be worth doing by changing consequences.

5.1.2.7 Choosing the Right Condition-based Technology

The choice of technology for condition-based maintenance will depend on several factors, including the type of asset being monitored, the complexity of the system, the amount of available data, availability of people, the warning period, accessibility, costs, and the required level of accuracy. In the right environment, the use of the human senses still has its place in modern maintenance.

5.1.2.8 The P-F interval –when the group says 'I don't know'

One issue that sometimes comes up when facilitators seek to determine the interval of a condition-based task is that when asked 'what is the P-F interval?' the group answers 'I don't know'.

There are several techniques to deal with this:

1. Ask a precise question

It is not reasonable to ask a vague or imprecise question, yet expect a precise answer. If the group is expected to give a reasonably precise answer, they should be asked a reasonably precise question. Throughout the RCM process, precision (including performance standards where possible) is sought. Performance standards in the function give performance standards in the Functional Failure which allow performance standards in the Potential Failure to be defined. For example, consider a pipe that fails due to corrosion. A Potential Failure to Functional Failure diagram is shown below. The normal wall thickness is 3mm and the pipe would fail at a wall thickness of 0mm. The potential failure – when the wall is failing – is between the two figures. The figure chosen must be something greater than 0mm (otherwise the pipe has failed). The exact figure chosen would depend on the accuracy of the technique used to measure wall thickness but would ideally be somewhere near the middle. For this example let's take a figure of 1mm.

Rather than 'what is the p-f interval?' a better, more precise question is now:

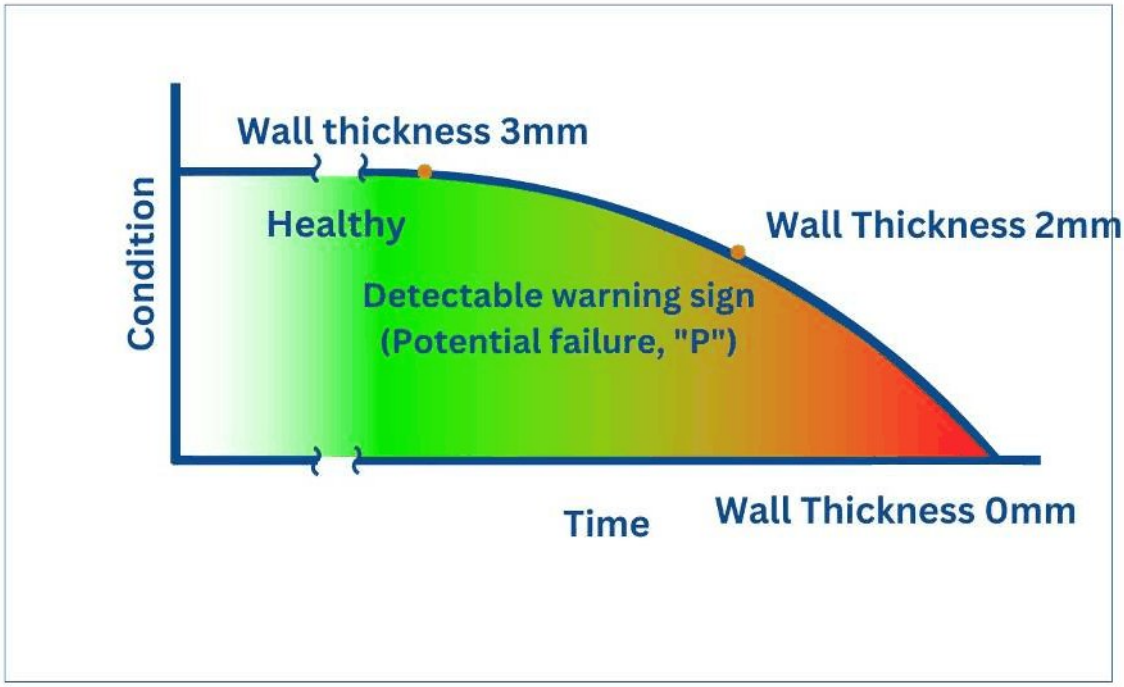

'what is the time from 1mm to 0mm?'

1. De)ne the failure mode properly

RCM deals with failures one at a time. - different failure modes are managed individually.

In this example, if the failure mode is normal corrosion, then a better question is:

'What is the time from a wall thickness of 1mm to a wall thickness of 0mm due to normal corrosion?'

(It is often worth defining deterioration as 'normal' (if applicable) to distinguish from abnormal corrosion which might happen quicker. However, these failure modes – if they exist – would be handled separately).

1. Stick to one failure mode at a time

It is often worth reminding the group at this stage that the P-F interval is being determined just for the failure mode under consideration. Sometimes the group may say:

…But what about sub-standard material? or …What about out of spec fluid?

….Or if someone were to hit the pipe in a fork lift truck? etc

These are, of course, separate failure modes that should be recorded separately (if applicable) and managed separately.

1. Ask the right people

The people best able to address this question are the people who have the best knowledge about the asset and how it behaves because they work with it every day.

CJ Other diagrams

ISSUE	RCA FEATURE	RCM FEATURE
Zero base	Takes existing equipment and strategies as the zero base	RCM takes no tasks as the zero base
Pattern F	No formal recognition of pattern F*	Formally recognizes pattern F
Technical feasibility of maintenance tasks	Technical feasibility criteria not defined*	Technical feasibility criteria defined
Worth doing criteria	Does not formally consider 'worth doing' criteria*	Formally considers 'worth doing' criteria
Standards	No international standards defining the process	Process defined by international standards
Tools	'Over 100 documented 'tools''	'Tools' built into the process

Definition / Explanation	What this means
• Proactive Tasks are undertaken before a failure occurs • Default Tasks deal with the failed state	• Proactive and reactive tasks are both maintenance tasks since maintenance includes actions to repair failures, and also to predict/ prevent/ detect failures • Failure-finding is categorised as default because it detects whether an item has failed • Redesign can sometimes be proactive and sometimes default
• Proactive maintenance is a cyclic activity in that it is carried out many times after fixed time periods (or measures of time)	• Physical redesign is not classified as a maintenance task because it is a 'one-off' activity rather than a cyclic/ recurring activity
• Planning defines the 'what', 'why' and 'how' • Scheduling refers to then timing of planned work	• Redesign is always planned (parts and tools identified and instructions defined ahead of time) and always scheduled (who is going to do and when they will do it is agreed) • Corrective work is never scheduled (because we don't know when it will occur). • Corrective work is sometimes planned (we might decide ahead of time to let the duty pump fail because it has a standby, but we have a plan about what to do if this happens). • Corrective work is sometimes not planned, for example if the consequence and likelihood of a failure are low we just respond afterwards.
• Predictive maintenance tasks look for potential failures (or warnings that the failure is occurring or about to occur), and the replacement or overhaul takes place based on the condition of the asset • Preventive maintenance tasks involve replacement or overhaul at a fixed measure of time regardless of the condition of the asset	• The categories of condition-based/ predictive tasks are the same as defined by John Moubray and include the use of the Human senses, Product quality monitoring, Performance Monitoring and Condition Monitoring • Condition Monitoring is the use of specialised equipment to monitor the condition of other equipment

The distinction between P-F Interval and Randomness: Example of a bearing

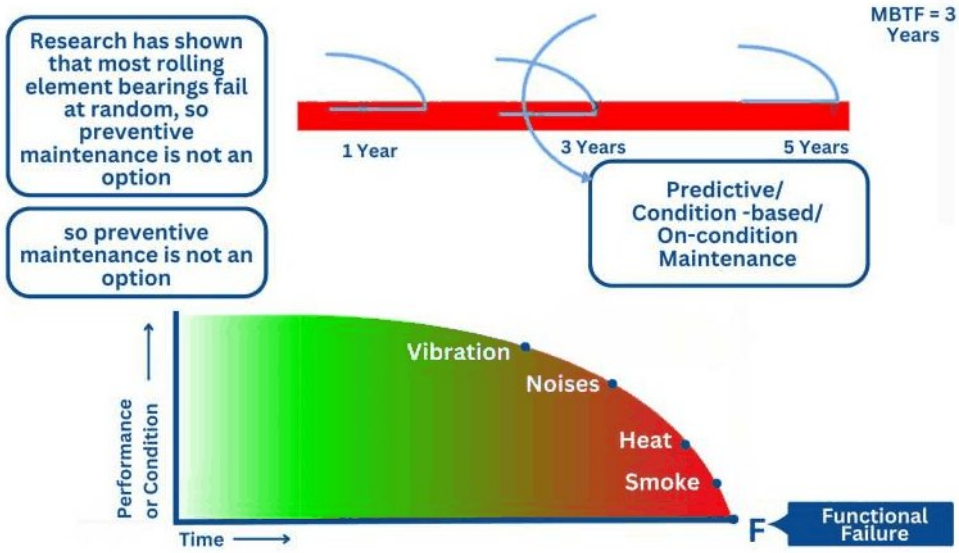

FEATURE	RCA FEATURE	RCM FEATURE
Failure Mode detail	Considers root cause of failure	Considers root cause of failure
Level of detail of failure mode	"5 why's" technique that involves repeatedly asking the question "why?" until the root cause is identified	Asks 'why' an appropriate number of times (until an appropriate failure management task is identified)
Types of failure considered	All causes (featured on fishbone diagram for prompt)	All causes (Deterioration. Design, Human Factors Operations)
Sources of information	Considers all sources of information – team based	Considers all sources of information – team based
Forward or backward looking?	Usually only looks at problems that have happened	Asks 'what could happen'? as well as 'what has happened'?
Targetted or comprehensive	Usually targeted at a particular problem	Usually carried out with particular objectives in mind. Can be targeted at particular problems

FEATURE	RCA FEATURE	RCM FEATURE
Trigger	Usually triggered by a problem or issue	Usually triggered by a poor performance/ bad actor report, or criticality analysis
Types of recommendations	Maintenance, physical redesign, training, procedures, contingency, spares, investigation	Maintenance, physical redesign, training, procedures, contingency, spares, investigation
Number of recommendations (analysis size)	Usually Five to twenty	Usually between ten and a few hundred
Maintenance task selection	No algorithm for selecting maintenance tasks in an appropriate order*	Algorithm for selecting maintenance tasks in an appropriate order (Decision Diagram)
Maintenance task selection interval	No rules for determining maintenance task interval*	Provides rules for determining maintenance task interval

5.1.3 Preventive Maintenance

Some Failure Modes are managed using Preventive Maintenance - the familiar concepts of overhaul or replacement. In RCM preventive tasks are considered if an on-condition task is not technically feasible and worth doing. Overhaul or Replacement actions are generally carried out while the assets / components are still serviceable (not yet failed). When we carry out overhaul or replacement actions we do not consider the condition we perform the task and that is that! These failures involve deterioration and increased risk of failure - the underlying pattern of failure is pattern B

Figure 29: The Life curve – Pattern B

Mechanisms of failure that may be age or number of cycles related include corrosion, erosion, wear, evaporation, fatigue and disassembly due to vibration. Examples of preventive maintenance on a car include oil replacement, oil and air filter replacements and cam/ timing belt replacements.

To assess the application of Overhauls/ Replacements (technical feasibility) we ask:

- Is there an age at which the likelihood of failure increases significantly and if so what is it?
- Do most last to this age?
- Can we return the capability to an acceptable level?

The distinction between Scheduled Overhaul and Scheduled Replacement

Overhauls usually involve the replacement of components of an item, whereas Scheduled replacements involve the replacement of the item as a whole. The distinction is really just one of levels, but the principles are the same. Therefore, the questions in the Decision

Diagram are the same for both types of task - overhauls, where components are restored to an "as new" condition and replacements, where the asset or component as a whole is replaced. The overhaul question comes first because it is usually cheaper (and more environmentally desirable) to overhaul something rather than replace it. (Although it is accepted that this is not always the case, and sometimes assets cannot easily be overhauled and a simple replacement is cheaper or less risky).

5.1.4 Worth Doing

To be selected a task must be technically feasible and worth doing. If a task has been identified that is technically feasible, the worth doing questions need to be addressed. These are the second set of questions in the task selection box of the decision diagram.

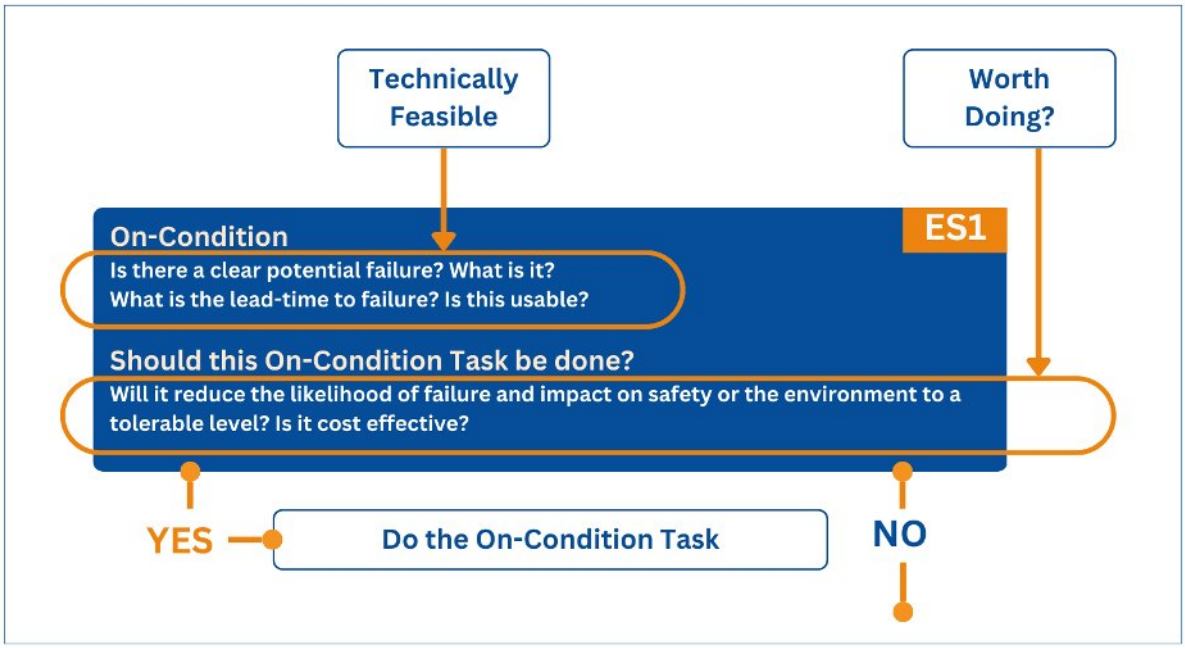

Figure 30: On-condition task evaluation – technically feasible and worth doing?

A task is worth doing if it reduces the consequences of failure to an extent that justifies the cost of doing the task. The "worth doing" questions therefore depend on the category of consequence of the failure.

5.1.4.1 Safety/ Environmental Consequences

For Safety / Environmental consequences a task is worth doing if it reduces the probability of failure to a tolerable level. This is usually a judgment about the risk of doing the task in relation to the risk of not doing the task.

5.1.4.2 Financial Consequences

For Economic consequences a task is worth doing if the total cost of doing the task is less than the total cost of not doing the task. This is usually expressed as a cost per year. Costs of failure include the following:

- Loss of or reduced production/supply
- Quality or Performance problems
- Loss of contract / compensation
- Higher operational costs
- Repair Costs (resources, spares)
- Secondary Damage
- Fines or other Penalties

While the costs of doing the task include the following:

- Loss of or reduced production/supply
- Repair Costs (resources, spares)
- Cost of doing the task (resources)

It is possible to do a full cost/ benefit calculation of a maintenance task to compare the costs of doing the task versus the cost of not doing the task. This is occasionally done for contentious or expensive decisions although for most tasks the decision is clear cut and a simple judgement is made.

5.2 Order of tasks in the Decision Diagram

5.2.1 Why Condition-based Maintenance first?

The Decision Diagram favours condition-based maintenance ahead of preventive maintenance. The reasons for this are as follows:

- Condition-based maintenance can be used for a larger number of Failure Modes. For example, most bearings fail at random, so preventive maintenance is not technically feasible. However, condition-based maintenance most certainly is technically feasible for most bearings. Indeed, there are usually numerous condition-based task options for bearings.
- Condition-based maintenance can often be carried out while the equipment is running. For example, thermography can usually be carried out while running (indeed it may be a requirement for the equipment to be running to carry out thermography)
- Condition-based maintenance is usually less invasive. Thermography and vibration analysis are good examples.
- Condition-based maintenance is often cheaper because it maximizes the life of components. The wear of a car tyre is a good example. This is a failure mode that has age-related characteristics and fixed time replacement (replacing the tyres at xx thousand miles) is technically feasible. However, it is cheaper to inspect the car tyres for wear and replace when necessary. This is because the tyres wear at a different rate and a safe policy of fixed time replacement would need to be conservative. Unfortunately, this would mean that the tyres that wear more slowly would be replaced when they still have some tread

left. A condition-based approach on the other hand, would ensure the tyres are replaced when they need to be replaced, and not before.

5.2.2 Why Maintenance First?

The Decision Diagram favours maintenance ahead of redesign. The reasons for this are as follows:

- Maintenance is often quicker
- Maintenance is often cheaper
- If redesigns involve adding hardware they often add more Functional Failures and therefore more maintenance
- Many organizations do not have significant resources available to implement large numbers of redesigns
- There is always an element of risk when changing something.

5.2.3 Proactive Maintenance ahead of Failure-Finding for hidden failures

A common misconception is that all hidden failures need to be managed by failure-finding. The Decision Diagram considers condition based, scheduled restoration and scheduled replacement ahead of failure finding. The reason for this is that if these strategies can be carried out, they are usually safer. This is because they can prevent the protective device every getting into a failed state, whereas, because failure finding involves checking to see if an item has failed, it will inevitably spend some time in the failed state. For example, even if a smoke detector is tested every hour, there is a chance that it will fail between each test, and therefore the user will be exposed to that risk. Decreasing the test interval will reduce the probability that it will be failed, but not eliminate it altogether.

Note however, that condition based, scheduled restoration and scheduled replacement are seldom applicable for protective devices because protective devices usually do not have wearout characteristics, and they do not exhibit potential failures.

5.2.4 Situations where a lower order task is clearly more effective

The task options in the Decision Diagram are presented in a logical and sensible order. However, there may be times when a lower order task (ie one that comes up later in the Decision Diagram) is preferable, even though the higher order task is technically feasible and worth doing. Two ways this can happen are:

1. Fixed time overhaul or replacement are preferred to condition based 2. Fixed time replacement is preferred to fixed time overhaul

Some examples of the first category include:

> 1. Small, low-cost items: For items that are cheap and easy to replace, such as seals or filters. It is often more cost-effective and less time-consuming to simply replace them at fixed intervals rather than constantly monitoring their condition. The same applies to the oil in our cars. Compared to a policy of No Scheduled Maintenance, it is surely technically feasible and worth doing to take a sample of oil, send it to the lab, check for oil breakdown and contaminants and replace when necessary. However, given the small volume of oil and its relatively low price, it is cheaper to replace the oil at a fixed interval.

2. Items with a well-understood life expectancy: Examples are often the same as above and might include some small batteries.
3. Equipment exposed to harsh or inaccessible environments: Equipment that is exposed to harsh conditions like high temperatures, high pressure, or corrosive environments may degrade in a manner that is difficult to monitor. In these cases, scheduled replacement can ensure that the equipment is always in good condition. Examples could include certain components in industrial, mining, or offshore operations. Similarly, some components may be difficult to access.
4. Large interventions – Large maintenance interventions, often referred to as shutdowns, turnarounds, or outages, are usually carried out at fixed intervals in large, complex systems such as a refinery hydrocracker or a warship. A major reason for setting a fixed time is planning - by scheduling these large maintenance activities at fixed times, it allows for better planning and coordination with other parts of the operation and teams of people (including external contractors) can be organised in good time. Also, production schedules can be adjusted to accommodate the downtime, and contingency plans can be put in place

Examples where fixed time replacement is preferred to fixed time overhaul include:

1. Low-cost or consumable components: In scenarios where components are relatively inexpensive or are designed to be replaced rather than repaired, like certain types of filters, or seals, replacement is often more cost-effective and less time-consuming than an overhaul.
2. Components with limited repair potential: Some components may have limited potential for effective repair or refurbishment due to their design or the nature of wear they experience.
3. When overhaul costs exceed replacement costs: In cases where the cost of overhauling a component exceeds the cost of a new component, replacement becomes the better option. This might occur with certain types of pumps, motors, or other mechanical devices

5.2.5 Measures of time

Note that maintenance tasks do not necessarily need to be based on calendar time. Other measures of time can be used. Examples include:

1. Operational Hours: This is the time that a piece of equipment or system has been in operation. This measure is often used for heavy machinery or vehicles where maintenance tasks are scheduled based on hours of operation. For example, an aircraft engine might require maintenance after a certain number of flight hours.
2. Cycles: In some contexts, maintenance is scheduled based on the number of cycles a piece of equipment has completed. A cycle represents the completion of a start-to-finish operational process. For instance, in the aviation industry, an airplane might require certain maintenance tasks after a specified number of takeoff and landing cycles.

Composite or alternative measures of time

Some tasks may not be based on a single measure of time. A good example is (again!) oil changes on a car. For example, a car manufacturer might recommend changing a car's oil every 5,000 to 7,500 miles or every six months, whichever comes first.

This is based on normal driving conditions, which means that if you frequently drive in harsh conditions, you may need to more frequent oil changes. Consequently some cars recommend oil changes based on an algorithm. These may base the oil change recommendation on actual data from the vehicle. At the heart of the equation are still mileage, revolu-

tions and hours of operation, but oil and coolant temperatures are also taken into account to see if the engine has had time to burn off condensates in the oil.

Another example of a composite measure from another industry is downhole drilling tools. Maintenance tasks might be specified. ‘After every job or every 1000 running hours or after every high shock’

CHAPTER 6

Default Tasks

The 7 master questions of RCM are represented below. This chapter addresses the seventh of these – what should be done if prediction or prevention is not possible? so?

1. What do we want the asset to do? – Functions
2. How can it fail? – Functional Failures
3. What causes the Functional Failures? – Failure Modes
4. What happens when a failure occurs? – Failure Effects
5. How much does each failure matter? – Consequences
6. Can we predict or prevent failure and should we be doing so? – Proactive Tasks
7. How should we manage the failure if prediction or prevention is not an option?

– Default Tasks

In the previous chapter we considered what happens down each of the columns of the Decision Diagram as an appropriate failure management policy is sought. In this chapter we will consider the various options at the bottom of the Decision Diagram, if a proactive maintenance task is not possible.

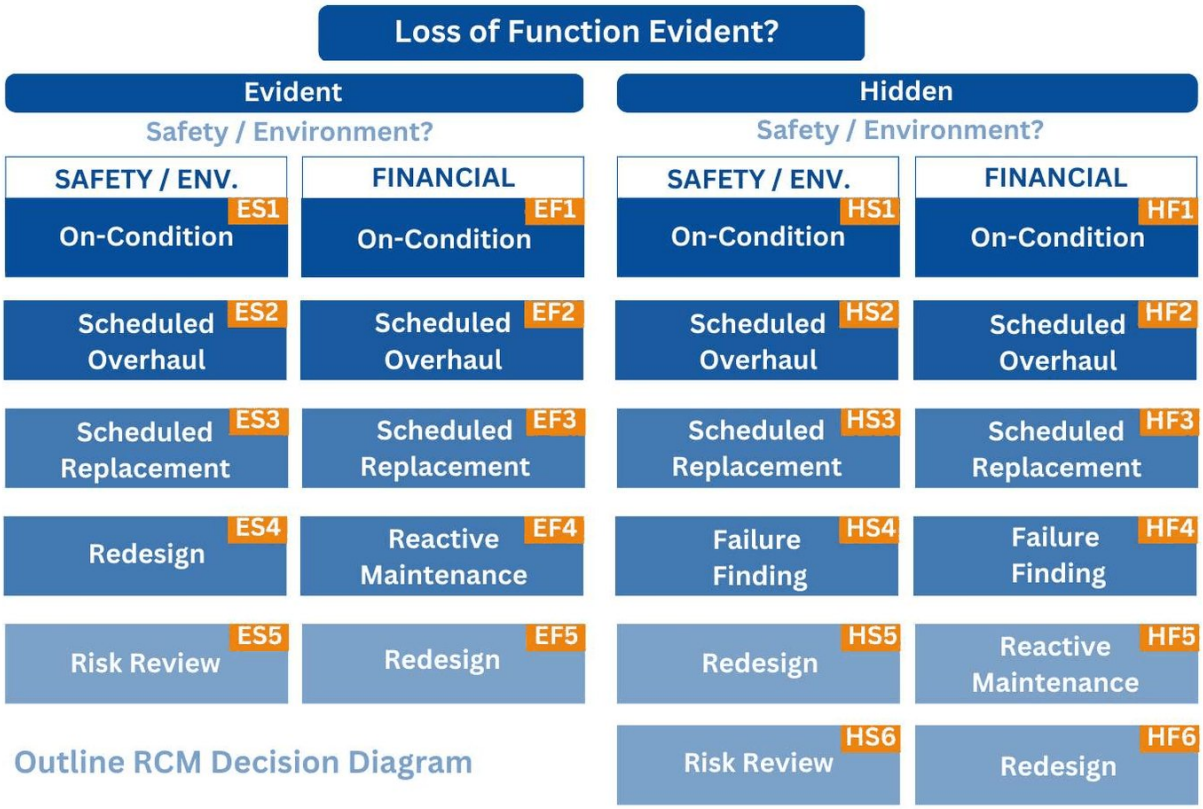

Figure 30: Outline RCM Decision Diagram

The options that are considered in this chapter are Redesign, Reactive Maintenance, Risk Management and Failure Finding.

6.1 Redesign

If On-condition or Preventive maintenance are not appropriate we may consider Redesign. In RCM the term 'redesign' has a broad meaning and may mean:

- Physical Modification
- Change or reinforcement of Operating Procedure
- Change what we do by Training
- Improve Response Time
- Modify Spares holding (to reduce downtime)

In fact, the term 'one-off change' or 'one-time change' is probably more apt (as opposed to maintenance which is a cyclic activity). Redesign suggestions also need to be worth doing and cost effective. During the RCM analysis, redesigns are not evaluated in detail. If a redesign suggestion looks attractive it should be recorded and the team should move on. However, the team should bear the following questions in mind when suggesting redesigns because redesign suggestions need to be kept manageable.

Safety/Environmental Consequences	Financial Consequences
Will it reduce the likelihood of failure to a tolerable level?	How much will it reduce the cost of failure (Lost production, Quality Issues, Secondary damage, Repair actions, Fines?)
Are the timescales and costs of design, implementation, certification etc. tolerable?	Is the investment worthwhile considering the timescales and costs of design, implementation, certification etc .
Is it the most cost-effective option?	

Figure 31: Redesign questions

A good analysis usually has a decent number of redesign proposals. In many ways this shows that the analysis is a 'real' analysis because these sorts of failure modes are often specific to the asset in question and it shows the team has analysed their asset. Deterioration failure modes on the other hand are more generic. The following diagrams show examples of the large number of redesigns that can sometimes come up in an analysis.

Selected DWD Pump failure management policies

TYPE OF ISSUE	RECOMMENDATION
Poor maintenance procedure	Ensure seal water bleed procedure is properly defined
Maintenance competence	Ensure all relevant maintenance personnel are properly trained in the use of the specific laser alignment equipment that is available on site
Procurement issues	Ensure seal fluid supplied to tank is clean
Poor quality components	Ensure Journal bearings procured from reputable source
Operator error	Follow procedure to prevent reverse pressurisation of mechanical seal during shutdown
Maintenance error	Ensure correct orientation of the thrust bearing pads during installation
Spares issues	Ensure filter spring is in SAP and hold spare spring in stores
Logistics issues	Improve local cranes to avoid delays caused by need to obtain cherry picker

TYPE OF ISSUE	RECOMMENDATION
Poor design causing unreliability	Replace seal water pump motor with 3 phase design
Poor maintainability	Consider elevating the seal water coolers to improve bleeding performance
Poor maintenance practices	Ensure measurements of distance between the two hubs - pump and motor - are taken as a reference and recorded, when the hubs are removed.
Poor housekeeping	Improve oil handling and storage
Poor operating practices	Do not leave duplex filter handle in the middle
Poor design affecting operability	Consider building a stand for operating manual suction valve
Poor design causing risk	Fit permanent mark on site to indicate what the settings should be for the nitrogen regulator, and ensuring measurement units are consistent
Poor quality repair	Ensure shrink disks are in good condition if re-used
Reduced output due to excessive cartridge wear	Consider Sume coated cartridges

Design, Training/Procedure/Competence, Can rising above want

KPI	COMMENTS
Maintenance Cost	Often total maintenance costs per year divided by replacement value of the facility, unit of output or manufacturing cost
Availability	May consider scheduled and/or unscheduled downtime
MTBF	Covered in availability but sometimes measured in its own right
Safety/ Environment/ Health	Extent to which the company complies with SHE laws and regulations. Includes LTIs
Preventive maintenance costs versus total maintenance costs	A measure of the extent to which maintenance is under (or out of) control. Does not reflect the fact that NSM may be a valid decision
Work orders completion	No of WOs with completion date less than or equal to scheduled completion date divided by number of WOs. There may be many versions of this measure, such as backlog
Maintenance Productivity	Total number of technician hrs entered for equipment related WOs divided by the total hrs per annum
Inventory value of spare parts versus asset replacement value	Measure of spare parts held. Needs to take context into account context.
Outsourced maintenance costs versus total maintenance costs	Costs of outsourced maintenance divided by the total maintenance cost. Needs to take context into account context.
Overtime percentage	A measure of the extent maintenance resource is adequate

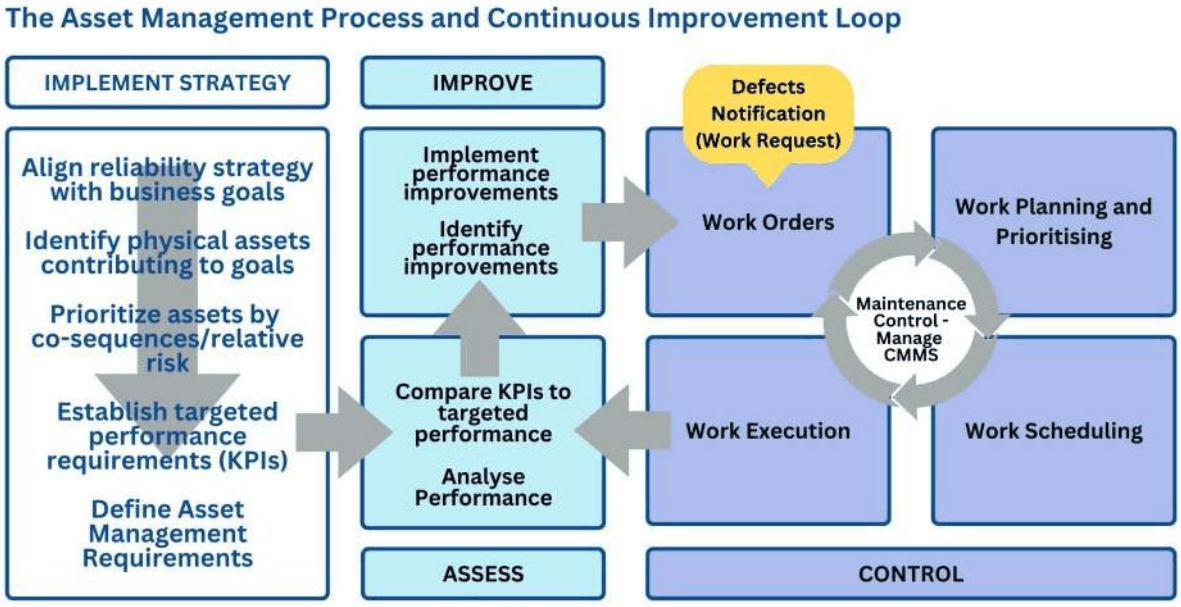

6.2 Risk Management

For failures with safety or environmental consequences, if we cannot find a suitable maintenance task and no Redesigns seem attractive, we also have the option of Risk Management. In certain cases when we have a failure with safety or environmental consequences (hidden or evident) the only option may be to live with the risk. If so, we need to be certain that all regulatory requirements are complied with and that the level of risk we face is tolerable, so the issue may need to be forwarded to the responsible persons for action, consideration and/or discussion.

6.3 Reactive Maintenance

If predictive or preventive maintenance is not an option we could "run to failure". In this case we simply let the failure occur and then repair it.

This can only be used if the failure has financial consequences.

"Reactive Maintenance Only" may also be called "Run-to-failure" or "No scheduled maintenance (NSM)".

Note that such a decision does not mean that the analyst has 'failed' – a large proportion of failures on aircraft are managed using No scheduled maintenance because of the high levels of redundancy and protection – the most appropriate policies depend on the equipment and the context.

6.4 Failure Finding

In the Failure Consequences module, we noted some "protective" devices whose failure was hidden. These types of devices include:

- Bursting disks
- Circuit breakers
- Overload trips

- Pressure relief valves
- Residual current devices

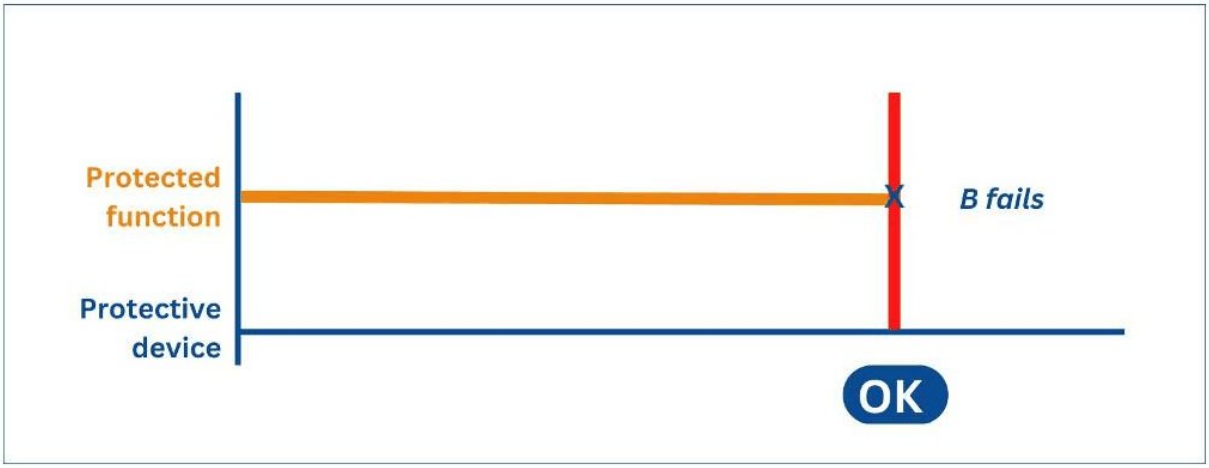

- Smoke detectors
- Standby pumps

Protective devices act to reduce or eliminate the consequences of failure of something else. This is the demand on the protective device.

Figure 31: As long as the protective device works when there is a demand the system is OK

For example if a standby pump works when it is called upon to do so, operation can continue. And if there is an overpressure and the relief valve works, the consequences are reduced.

However, if there is a demand while the protective device is in a failed state, a 'multiple failure' occurs. This is possible since all physical assets can fail.

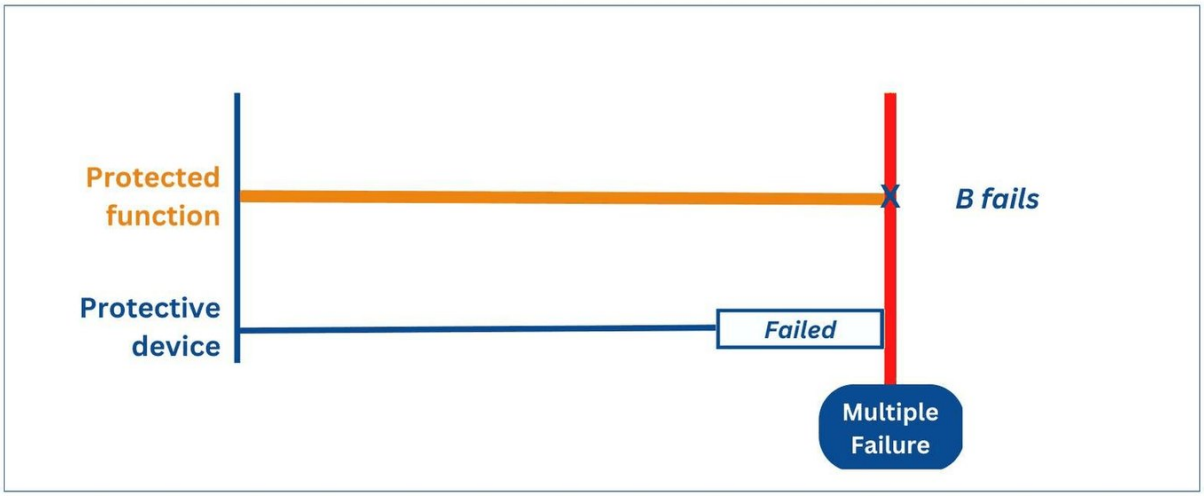

Figure 32: Multiple failure

6.4.1 The probability of a Multiple Failure

The probability of multiple failures and the associated management is the most mathematical part of RCM (however, the maths is not too onerous). Since the multiple failure is a result

of two independent events, the probability of a multiple failure is a product of two probabilities:

- The probability that there will be a demand on the protective device because of the failure of something else.
- The probability that the protective device will be in a failed state when there is a demand which is given by its unavailability (which is usually expressed as a percentage).

For example, if the demand rate on the protective device is once every 6 years, there is a one in six chance that the protective device will be needed per year.

And if the protective device has an availability of 80%, then it's unavailability is 20%, so there is a one in five chance that it will be in a failed state when called up on to operate. This means that the probability of the multiple failure is 1 in 6 x 1 in 5 = 1 in 30 years (or 1 in 30 per year).

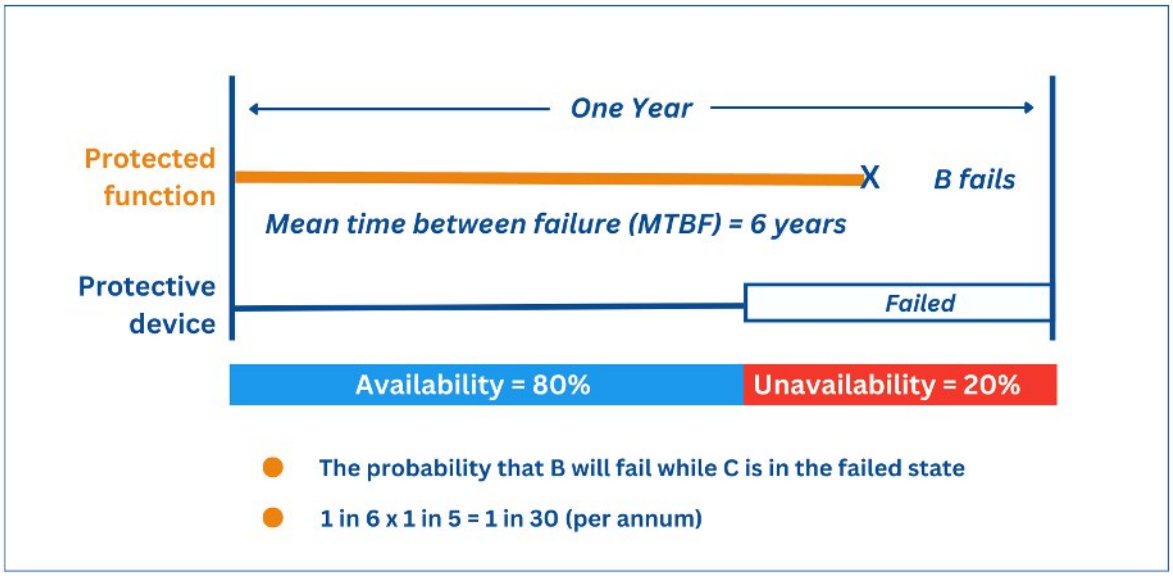

Figure 33: Calculating the probability of a multiple failure

When establishing a failure management policy for a hidden function, the first step is to decide what probability we are prepared to tolerate for the multiple failure.

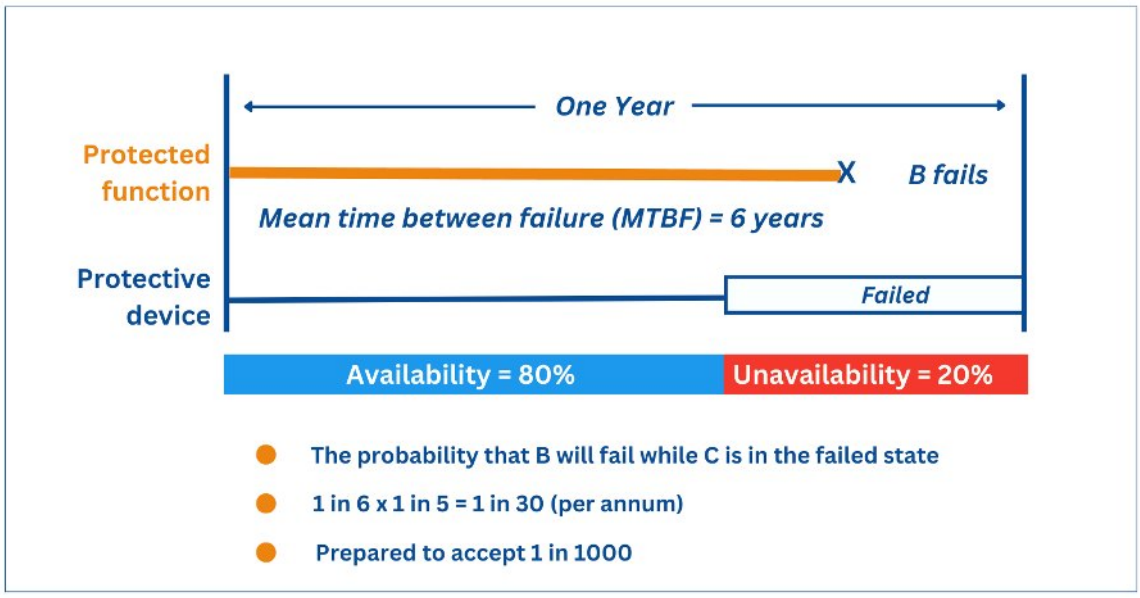

And then decide what has to be done to get it.

Figure 34: First step in managing multiple failure

There are many ways of reducing the probability of a multiple failure – through design and maintenance, both to the protective device and to the protected function. These include:

- By reducing the probability of failure of the protected function (by applying a suitable failure management policy or a change of design)
- And/or by increasing the availability of the protective device by:
- Preventing the failure of the protective device, or
- Periodically checking whether the protective device is still working and repairing it if it has failed (Failure Finding), or
- Modifying the system in some way

If appropriate, On-condition and Preventive methods may be used to manage Hidden Failures. The same technical questions as we asked earlier are applied to assess lead-time to failure or age and these tasks have been covered elsewhere in this document. However these can usually only be used for a minority of Hidden Failures.

The remainder of this chapter covers Failure Finding.

6.5 Failure Finding

Failure Finding is a maintenance option that only applies to Hidden Failures. We check whether the protective device / system is already failed (it either works or it doesn't). This is about managing the probability of the multiple failure by changing the availability of the protective device.

Whenever we use Failure Finding we should:

- Check the function of the entire protective system
- Be un-intrusive (minimum disturbance of the system)
- Not significantly increase the risk of a multiple failure

6.5.1 The frequency of a Failure Finding task – simple formula

How often we do Failure Finding depends on:

- The reliability of the protective device
- its mean-time-between-failure or "MTBF"
- The availability we require, which is driven by the consequences of failure

A simple formula for the failure finding interval using these parameters is given by:

2 x MTBF x (100% - A%) / 100

Where A = availability (expressed as a percentage).

The availability we require is dependent on the consequences of the multiple failure:

- for non-safety incidents we may accept availabilities of say 95–98%
- for safety related incidents we may want availabilities of say 99 – 99.9% or even greater device
- 95% availability means 5% unavailability, or a 1 in 20 chance it won't work when we want
- 99% availability means 1% unavailability, or a 1 in 100 chance it won't work when we want
- 99.9% availability means 0.1% unavailability, or a 1 in 1,000 chance it will not work when we want

Note that more rigorous equations are available including one that derives availability more scientifically, but these are beyond the scope of this document.

6.5.2 How to determine MTBF

To determine the MTBF of the organisation's assets we may have some good history records, alternatively national or shared records may be available, such as:

- NaFIRS (Electrical Transmission and Distribution)
- OREDA (Oil and Gas)
- Reliability, Maintainability and Risk – a book by DJ Smith
- Sometimes we have to rely on our own experience and observation (which is often the best source)

6.5.3 Failure Finding in the Decision Diagram

To assess the application of Failure Finding we ask:

- Is there a way to functionally check the protective device or system as a whole? – What is it?
- Can the check be done safely?
- How often must this be done?
- Is this practical and cost effective?

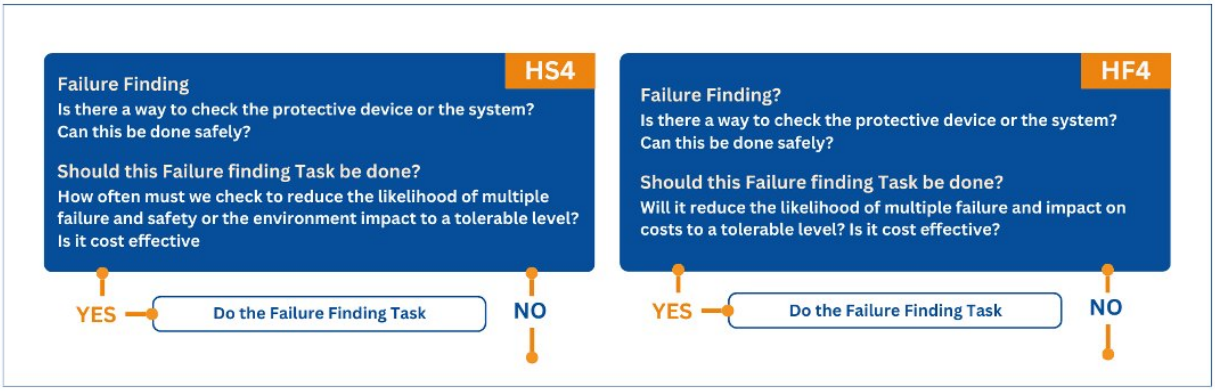

Figure 35 Decision Diagram Failure Finding boxes

6.5.4 Failure Finding – Other Formulae

The equation given above is a simple and less rigorous method of determining the interval of a failure-finding task. There are alternative, more rigorous methods. These are covered in this section. But first, we will consider what determines the interval of a failure finding task.

What determines the interval of a failure-)nding task?

Protective devices are not all equal! Some are more reliable than others, some are a different rate of demand than others, and some have higher consequences if they fail and they are needed and a multiple failure results. Therefore (provided they can be tested) the interval of a failure-finding task should not be the same for all protective devices. So, in general terms, the factors that affect how often a protective device should be tested are:

1. How often it fails
2. The rate of demand/ how often it is needed
3. The consequences of the multiple failure

If the consequences of the multiple failure are purely economic (ie it does not affect safety or the environment), money is spent carrying out the failure-finding task in order to save money by reducing the rate of multiple failures. In other words, we are spending money to save money. Unfortunately, this has a limit – there comes a point where spending more money is not worth it in terms of the money saved. There is therefore another factor that influences the frequency of the failure-finding task:

1. The cost of carrying out the task

And a financial optimisation could be carried out. There are three formulae for determining the failure-finding interval of a protective device. These are described below, and when to use them.

Availability based A less rigorous approach. Can be used for protective devices where the consequences of the multiple failure do not affect safety or the environment and do not have high economic consequences

Economic An economic optimisation. Can only be used for protective devices where the consequences of the multiple failure have economic consequences (ie they do not affect safety or the environment Risk based A rigorous approach that can be used for any failure finding interval calculation and must be used where the consequences of the multiple

failure affect safety or the environment. Note that this takes a figure for tolerable risk of the multiuple failure. The principles discussed in section xx CJ can be used for this purpose.

More complex protective systems

The discussion so far has focussed on single protective devices. Some systems are more complex:

- Parallel Systems – an example might be where there are two temperature sensors and the system shuts down if either of the sensors sends a certain signal (also known as a 2oo2 (2 out of 2) system). For example, sometimes relief valves are fitted in parallel
- Voting systems – an example might be where there are three temperature sensors and the system shuts down if any two of the three sensors sends a certain signal (also known as a 2oo3 (2 out of 3) system). Voting systems are common in, for example, fire alarm systems. The idea of a voting system is to be resilient against: o Failure of a sensor, because there are other sensors as backup o Spurious signals, because two sensors must send the signal for the trip to occur

Parallel and voting systems are used for critical protective devices with high consequences of the multiple failure and/or spurious shutdowns.

The mathematics behind the formulae to determine these task intervals is inevitably more complex. For all three formulae described above, there are variants for these more complex scenarios:

1. Parallel – Variations on each formula for protective devices using n parallel devices
2. Voting - Formula for protective devices using voting systems (k out of n devices (where n parallel devices are fitted and k devices need to send a signal for the alarm/trip to be actuated))

Recommended reference on Failure-Finding

Dr Mark Horton is probably the world authority about failure finding. He has written a book that covers this subject in more detail, including the derivation of the task intervals. (CJ Ref)

The following formulae are reproduced from his book with his permission.

CJ Formulae

Mathematical checks

For the formulae to be valid, certain mathematical requirements must be met.

1. The availability of the protective device must be at least 90% (ideally 95%). This is a requirement for the formulae to be valid mathematically
2. The demand rate for the protective device must be significantly less frequent than the failure-finding task interval. For example, there is no point testing a device every year if the demand is every 3 months, because the device is being tested by the system.

Both subjects are covered in Dr Horton's book.

The formulae can be built into an Excel spreadsheet and the calculator is available in the Reliability Management RCM App. The mathematical checks can be built ion to the logic, and flag a warning if any of the 'rules' have been breached.

More practical aspects of failure)nding

When considering failure finding, the most important question in the decision diagram is the first one; Is it possible to check if the item has failed? There is no point optimising the interval of a task that is not possible (or is risky) in the first place.

Data sources

Exact figures are usually not available and assumptions have to be made. Failure- Finding formulae enable us to be scientific about how we use our estimates. This is better than all the alternatives.

Understanding the system

With complex systems, most value is usually gained from understanding how they work and therefore how the failure-finding formulae are to be applied.

6.6 Categorisation of Failure Management Tasks

The table, and explanations below, show how many of the widely used terms in RCM and maintenance can be categorised. When using such terms it is important to understand what is meant, and if discussing with someone else, that their understanding is the same. These categories are covered in this book, not to definitively state what is the right and wrong meaning of particular words, but to ensure clarity about what is meant in this book, in case the reader has different interpretations.

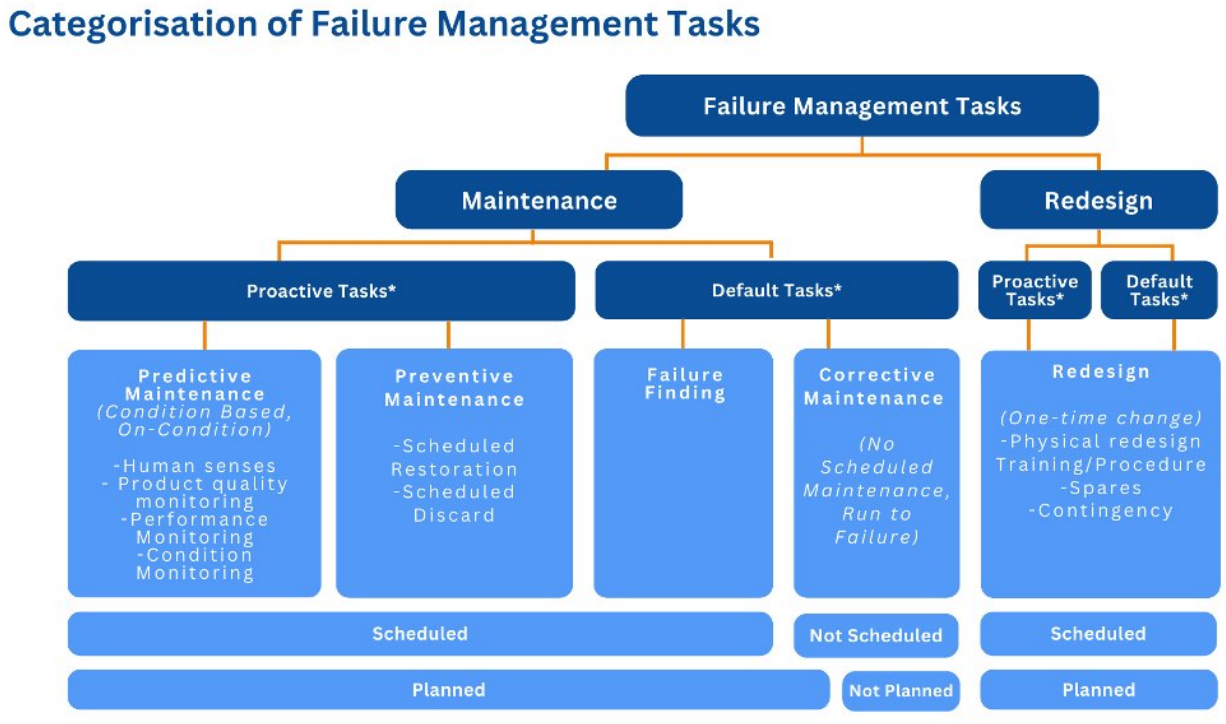

6.7 Basis of task intervals - a summary

RCM provides the means to determine which failure management tasks should be carried out. It also provides the technical basis for maintenance task intervals. This is summarised in the table below.

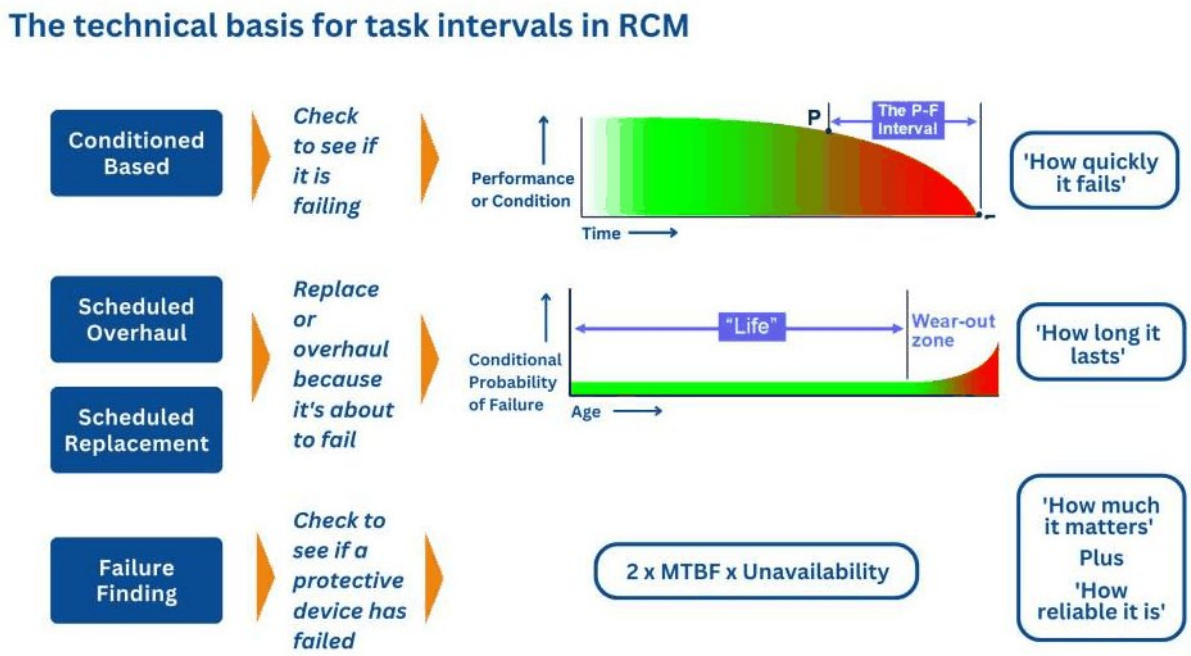

CHAPTER 7

Implementing and Applying RCM

This chapter deals with implementing RCM – ensuring the RCM decisions and recommendations actually happen, and applying RCM – setting up and selecting analyses that will be successful. However, it starts with a consideration of the alternatives to RCM for designing an asset management programme.

7.1 Alternatives to RCM

RCM is a systematic, structured, logical, pragmatic and scientific approach to developing a maintenance and asset management programme. It is worth considering the alternatives ie what other approaches can be used to determine the maintenance programme. These alternatives are discussed below. If the alternatives are considered it is clear that these all have massive drawbacks.

7.1.1 Doing Nothing

This approach cannot be a serious consideration for any asset with high consequences. Most importantly, such an approach may pose safety risks to personnel. Unplanned failures can also lead to unplanned downtime, with high costs of production loss, and high repair costs. and negatively impact productivity. Reactive maintenance often involves costly repairs or replacements, as opposed to preventive or predictive maintenance, which can minimize costs by addressing issues before they become critical. This leads to the question; how to determine what maintenance needs to be done, and at what interval? The answer of course, is RCM.

7.1.2 Doing what we've always done

This approach lacks adaptability and does not consider changes in the operating environment, changes to the equipment, or other factors that may affect the equipment's performance. Furthermore, doing what you've always done with mean that you will get what you have always got, and what if this is not acceptable? This approach can therefore lead to unnecessary maintenance, wasted resources, and missed opportunities for improvement. By sticking to the status quo, organizations may overlook opportunities to optimize maintenance tasks, streamline operations, and reduce costs. However, some traditional maintenance tasks may still be valid, so if all existing maintenance is rejected, we may be throwing out good tasks with bad ones. This leads to the question; how to assess which existing tasks are needed and which are not, and what process should we use to identify new tasks? The answer of course, is RCM.

7.1.3 Manufacturers Recommendations

Many organisations use manufacturers' recommendations for their maintenance programme. A problem here is that manufacturers' recommendations are often based on generalized or worst-case assumptions and may not consider the specific use-case, operating environment, or unique factors affecting the equipment. Furthermore, manufacturers are

often not aware of the maintenance techniques available to the user (particularly condition-based techniques) and sometimes not aware of any instruments fitted to the wider process which can be used for condition-based tasks. Consequently, manufacturers recommendations may favour preventive maintenance, or invasive predictive tasks. Another point is that many manufacturers only consider their equipment in isolation and do not consider (and to be fair, are not aware of), the impact of their recommendations on the wider process. While many manufecturers recommendations are undoubtably valid, following manufacturers recommendations without challenge can therefore result in sub-optimal maintenance, increased costs, and reduced equipment lifespan. This leads to the question; how to assess which manufacturers recommended tasks are needed and which are not, what process should we use to identify new or better tasks? The answer of course, is RCM.

While many believe in adhering strictly to manufacturers' maintenance recommendations, it's crucial to recognize the value of questioning these guidelines. The reasons for this include:

1. Profit Motive: Some manufacturers may have an interest in selling spare parts, which could influence their recommended maintenance schedules.
2. Lack of Contextual Consideration: These recommendations are often based on worst-case scenarios and may not take into account the specific conditions under which the equipment is being used.
3. Ignorance of Condition-Based Techniques: Manufacturers may not be privy to the condition-based maintenance techniques available to users, which could be more effective and economical.
4. Insufficient Time or Experience: Those who design and specify equipment might not have the necessary time or experience to carefully consider all maintenance requirements, leading to overly conservative recommendations.
5. Component versus System Perspective: Manufacturers often focus on individual components rather than the system as a whole. This can result in overlooking the broader impact of maintenance activities on overall system performance.
6. Unchallenged Sub-Contractor Recommendations: Suggestions from subcontractors can sometimes be incorporated into the manufacturer's maintenance recommendations without proper scrutiny.
7. Outdated Techniques: Maintenance methodologies continuously evolve. There may be newer, more efficient techniques available since the original equipment manufacturer (OEM) recommendations were written.

Hence, while manufacturers' recommendations provide a useful starting point, it's wise to adapt them to the specific usage conditions, available resources, and maintenance capabilities. Always consider these recommendations in conjunction with real-world experience and the latest best practices in maintenance

7.1.4 Adding a new maintenance task whenever there is a failure

Many companies do this, particularly for failures that have high consequences. Indeed,

if there is a failure, managers and customers may ask,		what maintenance did you do,

and what maintenance are you going to do?	The problem here is that tasks may be	
added without sufficient consideration of the root causes, and in many cases, the root		

causes may not be due to reasons that maintenance can address. This approach may then lead to an excessive and unnecessary number of maintenance tasks, increased costs, and decreased productivity. The principle of pattern F also means that there is a risk that unnecessary tasks may actually cause more failures. On the other hand, some critical failures do need to be addressed through maintenance. This leads to the question; how to assess which critical failures tasks need to be managed through maintenance tasks, how can we identify them proactively, and what process should we use to identify such tasks? The answer of course, is RCM.

7.1.5 Taking the equipment apart and replacing components whenever the equipment is accessible

It makes sense to do maintenance when the equipment is available because it is down for some other reason. However, it is counterproductive to do tasks that are not necessary just because the equipment is available. Unnecessary maintenance tasks can result in extended periods of downtime, higher maintenance costs and potential for damage because disassembling and reassembling equipment may introduce the risk of damage or errors. This leads to the question; how to identify what really needs to be done when the equipment is down, and what does not? The answer of course, is RCM.

The Reliability Centered Maintenance (RCM) approach stands out as the superior choice for setting a maintenance program due to its objective, systematic, and scientific methodology. Other alternatives pale in comparison for several reasons:

1. Doing Nothing: This approach can lead to severe consequences, particularly if a failure has high implications, including safety risks.
2. Reactive Maintenance: Merely adding a task following a failure is short-sighted. It overlooks the potential causes of failures that might be outside the maintenance team's control, and doesn't account for instances where failures are infrequent or consequences are negligible.
3. Adhering to Tradition: Sticking to what has always been done dismisses the possibility of change, or the need for improvement if the current approach hasn't been particularly effective.
4. Convenience-Based Maintenance: Performing maintenance simply because access is currently available may lead to unnecessary tasks, wasting resources on maintaining equipment that doesn't need it.
5. Research-Oriented Approach: Investigating why equipment is still working might be valuable for research purposes but doesn't directly contribute to the maintenance program.
6. Training-Based Approach: While it's essential for maintenance personnel to understand the equipment, using maintenance tasks for this purpose leads to inefficient use of resources.

7. Experience-Based Approach: Relying solely on experience without a systematic, scientific process that is documented might lead to valuable knowledge getting lost or overlooked.

In contrast, the RCM approach provides a systematic and scientific methodology for maintenance that considers the operational context of the equipment, the consequences of failure, and utilizes the knowledge and experience of the maintenance team in a structured way. It helps in setting a maintenance program that is effective, efficient, and adaptable to changes over time.

7.2 Implementing RCM

Implementation and change management are critical aspects of any project, especially those that involve significant changes to processes, systems, or organizational structures.

> 1. Implementation: This is the stage where planned actions, strategies, or systems are put into practice. Effective implementation ensures that the project delivers the expected outcomes and benefits. Without proper implementation, even the most well-planned project can fail. It involves coordinating activities, allocating resources, managing stakeholders, and monitoring progress. Effective implementation also requires consideration of practical aspects, like user training and support structures, to ensure that the project's outputs are effectively used.
> 2. Change Management: Projects often involve changes - to processes, technologies, job roles, or even organizational culture. Change management is the process of helping individuals, teams, and organizations transition from their current state to a desired future state. It's crucial because changes can be disruptive and met with resistance. Change management helps in reducing resistance, increasing engagement, and ensuring that changes are effectively adopted and benefits realized.
>
> Change management involves activities like communication of the need for change and its benefits, consultation with those affected, providing training and support, and managing resistance.
>
> Both implementation and change management are crucial for achieving project goals and ensuring that the benefits of a project are fully realized. They help in transitioning from plans to action and from current practices to new ways of working, while minimizing disruption and resistance.

The outcomes of RCM are:

- Maintenance schedules
- Operating procedures (high frequency maintenance tasks carried out by the operator)
- Suggested one-off changes including physical redesigns, changes or reinforcements of procedures, training recommendations and spares recommendations

If these outcomes are not implemented then the benefits will not be realized. Some of the steps for implementation are described in the sections below:

7.3 The Audit

Before implementation, managers with overall responsibility for the assets should satisfy themselves that all information and decisions are sensible and defensible. Areas of disagreement should be discussed with the people who performed the analysis

7.4 Task Descriptions

Describe routine tasks as fully as possible. Descriptions should state:

- Exactly what is to be done
- What performance standards (if any) apply to the task
- What part of the machine or which component is affected (If necessary, descriptions should include component or sub-assembly numbers).

The more thoroughly the tasks are described on the worksheet, the less time and effort needs to be spent later on procedure writing.

Schedule planning is simplified if schedule frequencies are multiples of each other so schedule frequencies in the RCM worksheet may be adjusted so that they are sensible multiples.

> Sure, here are some examples of poorly written maintenance tasks and their precise counterparts:
>
> Poorly Written:
>
> 1. "Check the pump." - This is unclear about what exactly should be checked on the pump and how the check should be performed.
> 2. "Fix the conveyor." - This does not specify what problem needs to be fixed or what steps should be taken.
> 3. "Inspect equipment." - This lacks specifics on which equipment to inspect, what to look for, and how to perform the inspection.
>
> Precisely Written:
>
> 1. "Inspect the pump for signs of leaks, unusual vibrations, or abnormal noises. Check the pressure and temperature readings against the manufacturer's specified optimal range."
> 2. "Troubleshoot the conveyor system for issues with belt alignment. Check the motor and rollers for signs of wear. If belt misalignment is observed, adjust according to the equipment manual's guidelines."
> 3. "Perform a visual inspection of the HVAC system, checking for any signs of corrosion, leaks, or damage to the ductwork and filters. Verify that the system's temperature and pressure readings fall within the manufacturer's specified parameters."
>
> A precisely written maintenance task includes specific steps to be taken, clear descriptions of what to look for, and references to standards or guidelines when appropriate. This enables the maintenance personnel to carry out the task accurately and effectively, ensuring the reliability and longevity of the equipment

7.4.1 Operating Procedures

Many high frequency maintenance tasks generated as a result of an RCM analysis are carried out operators. Tasks to be done by operators should be woven into existing Standard Operating Procedures (SOP's) wherever possible.

7.4.2 Redesign

"Redesign" entails upgrading the capability of a process so that it can do what its users want. This is done by making a one-time change to any of the following:

- the physical design of the asset
- the capability of the maintainer or operator
- maintenance or operating procedures

Implementing these changes will require the support and co-operation of the engineering/ projects function, the training department and operations management.

7.5 Applying RCM

It is clear that answering the 7 questions below - gathering the information in the information worksheet, and making sensible decisions using the decision worksheet, requires inputs from different perspectives. No single person knows all there is to know about the asset.

1. What do we want the asset to do? - Functions
2. How can it fail? - Functional Failures
3. What causes the Functional Failures? - Failure Modes
4. What happens when a failure occurs? - Failure Effects
5. How much does each failure matter? - Consequences
6. Can we predict or prevent failure and should we? - Proactive Tasks
7. How should we manage the failure if prediction or prevention is not an option? - Default Tasks

These questions can only be answered sensibly by design, maintenance and operations people working together. RCM provides these people with a common set of values and a common language so that they can answer these questions together.

Therefore RCM can only be properly applied by a team.

Figure 37: RCM team composition

The benefits of a team approach include Validity, Ownership, Empowerment, Team building, Efficiency of Review, Learning and improved understanding.

7.5.1 Facilitators

Facilitators are specialists in RCM. Their job is to ask the questions and record the answers – to be the owner of the RCM process and help the review group through the RCM decisions – but the decisions are made by the groups

> An RCM (Reliability Centered Maintenance) facilitator plays a pivotal role in the RCM process, ensuring its effective implementation and guiding the RCM team throughout. Here are the main responsibilities of an RCM facilitator:
>
> 1. Lead the RCM Process: The RCM facilitator guides the RCM team through the steps of the RCM process, from defining the system under study to developing and implementing the maintenance strategy.
> 2. Coordinate Team Activities: The facilitator coordinates the activities of the RCM team, which often includes personnel from different disciplines such as maintenance, operations, engineering, and safety. This includes scheduling meetings, setting agendas, and ensuring the team stays on task.
> 3. Promote Team Collaboration: The facilitator promotes a collaborative environment where all team members feel comfortable sharing their knowledge and experience. They ensure all voices are heard and manage any conflicts that may arise.

4. Knowledge Transfer: The facilitator helps the team understand the principles and methodology of RCM, providing training as needed.
5. Ensure Quality of the RCM Analysis: The facilitator ensures that the RCM analysis is comprehensive, accurate, and meets the standards and guidelines for RCM. This includes reviewing the failure modes, consequences, and proposed tasks for completeness and correctness.
6. Manage Documentation: The facilitator manages the documentation of the RCM process, ensuring that the information captured is accurate, complete, and organized in a manner that supports future maintenance planning and decision making.
7. Facilitate Decision Making: The facilitator helps the team in making decisions about the appropriate maintenance strategies, ensuring they are based on a thorough understanding of the failure modes and consequences.
8. Promote Continuous Improvement: The facilitator promotes a culture of continuous improvement, encouraging the team to learn from the RCM process and to look for opportunities to improve the effectiveness and efficiency of maintenance activities.
9. Liaise with Management: The facilitator communicates the progress, findings, and recommendations of the RCM process to management, and secures their support for implementing the proposed maintenance strategies.

By effectively fulfilling these responsibilities, the RCM facilitator plays a crucial role in ensuring the success of the RCM process and the resulting improvement in equipment reliability and maintenance efficiency.

7.5.2 Which Assets to Analyse first

Most organisations do not apply RCM everywhere, but they do so on selected assets where the cost of doing an analysis is justified by the costs of not doing an analysis – the costs and risks associated with present unreliability. Some people argue that because RCM is a significant investment to do everywhere they should not start. This is not a valid argument – the important question is not where to finish, but where to start! Two approaches are usually used to determine which assets to analyse first – Bad Actors and Criticality Analysis.

7.5.2.1 Bad Actors

Bad Actors are assets that perform badly in relation to various KPIs, such as availability, reliability, cost or safety or environmental performance. This approach is widely used to determine which assets to carry out RCM on.

A simple way of describing this is to ask which assets are causing the most ‘pain’.

A typical list of common KPIs is given below. The first 4 are the most common reasons for carrying out an RCM analysis.

KPI	COMMENTS
Maintenance Cost	Often total maintenance costs per year divided by replacement value of the facility, unit of output or manufacturing cost
Availability	May consider scheduled and/or unscheduled downtime
MTBF	Covered in availability but sometimes measured in its own right
Safety/ Environment/ Health	Extent to which the company complies with SHE laws and regulations. Includes LTIs
Preventive maintenance costs versus total maintenance costs	A measure of the extent to which maintenance is under (or out of) control. Does not reflect the fact that NSM may be a valid decision
Work orders completion	No of WOs with completion date less than or equal to scheduled completion date divided by number of WOs. There may be many versions of this measure, such as backlog
Maintenance Productivity	Total number of technician hrs entered for equipment related WOs divided by the total hrs per annum
Inventory value of spare parts versus asset replacement value	Measure of spare parts held. Needs to take context into account context.
Outsourced maintenance costs versus total maintenance costs	Costs of outsourced maintenance divided by the total maintenance cost. Needs to take context into account context.
Overtime percentage	A measure of the extent maintenance resource is adequate

7.5.2.2 Criticality Analysis

Asset criticality is widely used to select assets for an RCM analysis. (Note the distinction between Criticality of the Asset as a whole, and Criticality of individual Failure Modes (which is covered elsewhere in this book)).

The criticality of an asset can be determined by considering two factors: the likelihood of the failure occurring and the consequences of that failure. The likelihood of failure can be determined by considering historical data, generic databases. The consequences of failure can be determined by analyzing the impact of the failure on safety, the environment, and the organization's operations.

Once the likelihood and consequences of failure have been determined, each asset can be assigned a criticality level, often on a scale from 1 to 5, with 1 being the most critical and 5 being the least critical.

CONSEQUENCE		PROBABILITY				
	Definitions	**1**	**2**	**3**	**4**	**5**
		Remote Event hardly heard of. May occur but in exceptional circumstances MTBF 1000 years or greater	Unlikely Event occurs in the industry, and could on this site, but doubtful MTBF 100-1000 years	Possible Event occurs from time to time on this site. MTBF 10-100 years	Probable Event is expected to occur on this site MTBF 1-10 years	Very Likely History of failures exists from previous or similar designs. MTBF < 1 year.
E	Catastrophic, immediate equipment failures with major damage to equipment, property and life possible. Safety hazard. Plant shutdown greater than 1 month or equivalent partial shutdown	6	7	8	9	10
D	Plant shutdown 2 days to 1 month or equivalent reduced capacity. Collateral damage to major equipment incurred. Requires major overhauls to restart. High component and/or labor costs needed to repair the failure	5	6	7	8	9
C	Plant shutdown 2 hours to 2 days or equivalent reduced capacity.Possible minor collateral damage to equipment in plant. May require minor overhauls to restart. Moderate component and/or labor costs needed to repair the failure.	4	5	6	7	8
B	Plant shutdown up to 2 hours or equivalent reduced capacity. No collateral damage incurred. Low component and/or labor costs needed to repair the failure.	3	4	5	6	7
A	No effect on plant. Negligible component and/or labor costs needed to repair the failure.	2	3	4	5	6

Which approach is preferable – bad actor analysis or criticality analysis?

> Carrying out reliability improvement based on bad actor analysis focuses on equipment that causes frequent operational issues, downtime, and maintenance costs. These "bad actors" often consume a disproportionate amount of maintenance resources. By focusing on these problematic assets, significant improvements can be made in overall reliability and operational efficiency.
>
> Criticality analysis, on the other hand, identifies the assets that would have the most significant impact on operations if they were to fail, irrespective of their current performance or reliability. While this is crucial for risk management and prioritizing maintenance activities, it does not necessarily target those assets presently causing the most frequent or ongoing problems.
>
> By focusing on bad actors first, organizations can often achieve quicker wins and more immediate improvements in reliability. Furthermore, this approach tends to optimize resource allocation by focusing on improving equipment that is known to cause frequent issues, rather than spending resources on equipment that, while critical, may already be performing reliably.

7.5.3 A Living Programme

RCM was created in the aircraft industry to allow decisions to be reviewed and improved in the light of experience. This is one of the benefits of an auditable process. This means RCM can be part of a continuous improvement cycle for maintenance. In order to do this is is recommended to review the analysis every 6 months to 1 year, or update when important new information becomes available.

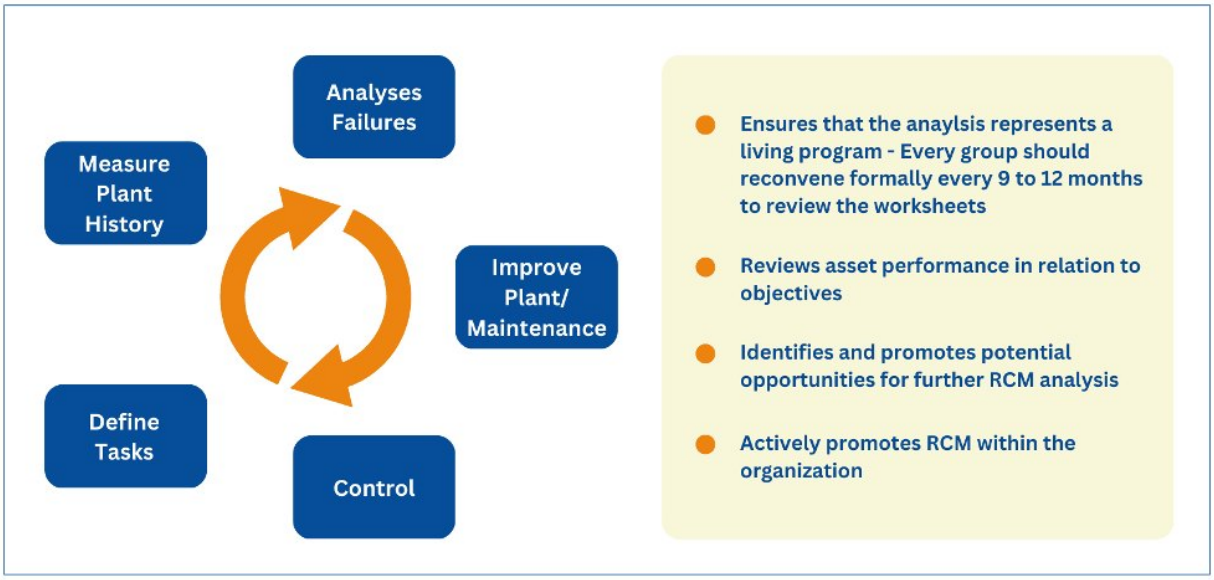

Figure 38: Continuous improvement using RCM

CJ better diagram here

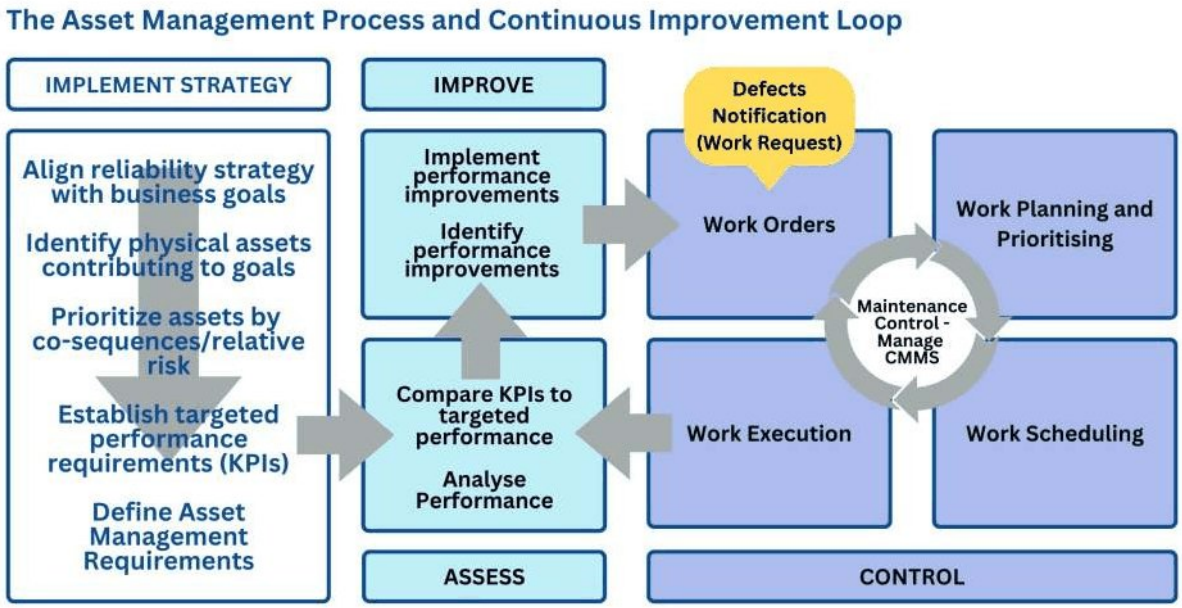

7.6 Levels of Analysis

RCM can be applied at different levels.

CJ different levels of training

7.7 Change Management and RCM implementation

The successful application of RCM would normally recommend more than a few adjustments to maintenance tasks. It should usually result in the adoption of new methodologies and practices.

Ensuring these are fully and properly adopted requires effective Change Management. Change management is the process of planning, implementing, and managing changes to an organization's processes, systems, and structures. If proper Change Management is not

carried out, the RCM implementation may fail because benefits will not be realised in practice. Some of the major challenges are discussed further below.

7.7.1 Resistance to Change

Many employees may be resistant to the idea of changing the way they have been doing things for years. They may be comfortable with the current maintenance practices and feel that RCM recommendations will disrupt their work. One example may be that maintenance planners are more comfortable with traditional Preventive Maintenance and feel that a change to a more condition-base strategy will make their work less predictable. It may be necessary to address this resistance by communicating the benefits of RCM and involving employees in the decision making and change process.

7.7.2 Alignment with Business Goals

RCM implementation is not just about maintenance. It is also about aligning maintenance practices with the organization's business goals. It is important to ensure that the implementation of RCM is aligned with the organization's overall strategic objectives, and that this alignment is communicated.

7.7.3 Training and Development

RCM requires a significant shift in mindset and skillset for maintenance and operations personnel. Employees may need to receive training and development to successfully implement RCM principles. This includes training on the RCM methodology itself, as well as on any new tools or technologies that may be identified in the new Failure Management programme.

7.7.4 Communication and Engagement

Effective communication and engagement are critical when implementing RCM. Employees need to understand why the change is necessary and how it will benefit the organization. It can also help identify and address any concerns or questions employees may have.

7.7.5 Continuous Improvement

RCM is a continuous improvement process. It can help ensure that the organization continues to evolve and improve its maintenance practices over time. This includes regularly reviewing and updating RCM processes and procedures, as well as soliciting feedback from employees on how to further improve the maintenance strategy. In conclusion, introducing RCM principles into an organization is not just about implementing a new maintenance strategy. It also requires effective change management to ensure successful adoption and alignment with the organization's strategic objectives. By addressing resistance to change, aligning with business goals, providing training and development, communicating effectively, and continuously improving, organizations can successfully implement RCM and optimize their maintenance practices.

7.8 A Change Management Checklist

The checklist below can be used to help ensure RCM is successfully implemented in an organization:

	A Change Management Checklist
☐	1. Define the scope of the initiative, including the assets, systems, and processes that will be covered.
☐	2. Consider the goals and objectives of the initiative, such as improving asset reliability or availability, reducing maintenance cost etc.
☐	3. Identify the stakeholders who will be impacted by the initiative and involve them in the planning process.
☐	4. Develop a communication plan to keep stakeholders informed about the progress of the initiative. This may include a series of short presentations to management and the wider workforce.
☐	5. Establish performance metrics to select bad performing assets for analysis, and the effectiveness of the change initiative. If necessary, develop a system for collecting and analyzing data on asset performance and maintenance activities.
☐	6. As an alternative to selecting bad performing assets, identify and prioritize critical assets for RCM analysis.
☐	7. Develop a detailed plan for implementing reliability centered maintenance, including timelines and resource requirements for training, analysis and implementation.
☐	8. Train employees who will be directly involved in the project on the principles and practices of reliability centered maintenance.
☐	9. Carry out the studies according to the plan ensuring that the team is given the time and support required to carry out the studies properly.
☐	10. Review the analyses. They should be checked for the correct application of the process, and reviewed by Managers with overall responsibility for the assets, to ensure they agree with the recommendations.
☐	11. Identify and address potential barriers to the successful implementation of the change initiative.

☐	12. Communicate the results of the change initiative to stakeholders and celebrate successes.
☐	13. Implement the Failure Management Strategies for the assets
☐	14. Monitor asset performance and adjust the Failure Management strategy as needed.
☐	15. Embed reliability centered maintenance into the organization's culture and ensure that it is sustained over time.
☐	16. Continuously monitor and improve the reliability centered maintenance process to ensure ongoing effectiveness.

Why does RCM work?

The lifecycle of an asset, from initial design through to decommissioning, is typically divided into several distinct stages. Here are the common steps involved:

1. Conceptualization: This is the initial phase where the need for a new asset is identified. The asset's specifications, performance requirements, expected lifespan, and target cost are defined during this stage.
2. Design: In this phase, the asset is designed to meet the requirements identified in the conceptualization stage. This could involve creating blueprints, models, or prototypes. Considerations might include not just performance, but also ease of manufacture, maintenance, safety, and environmental impact.
3. Manufacturing/Construction: Once the design is finalized, the asset is manufactured or constructed. This involves sourcing materials, assembly, and quality checks to ensure that the asset meets the design specifications.
4. Installation and Commissioning: The asset is installed at its operating location, and commissioning activities are carried out to ensure it is operational and meets the specified performance criteria.
5. Operation: This is the phase where the asset is put to use. It's the longest phase in the asset's lifecycle and can span many years or even decades.
6. Maintenance and Repair: Throughout the operational phase, regular maintenance and repair activities are carried out to keep the asset in good working condition and extend its lifespan. This can include preventive maintenance, predictive maintenance, inspections, and necessary repairs or replacements.
7. Upgrades and Modifications: As technology progresses or needs change, the asset may be upgraded or modified to improve its performance, extend its lifespan, or adapt it to new uses.
8. Decommissioning: When the asset is no longer needed or has reached the end of its useful life, it is decommissioned. This involves safely taking the asset out of service and may also include steps to decontaminate, disassemble, and dispose of or recycle the asset.
9. Disposal/Recycling: After decommissioning, the asset is either disposed of in an environmentally safe manner or recycled. Recycling can involve repurposing or reclaiming materials from the asset for use in other applications.

Throughout all these stages, it's important to consider factors like cost, safety, environmental impact, regulatory compliance, and the needs and goals of the organization. Lifecycle management strategies, like lifecycle costing or asset lifecycle management, can help in making informed decisions at each stage to optimize the value of the asset over its entire lifecycle

Asset design can often be suboptimal due to several reasons:

1. Insufficient Understanding of Requirements: If the requirements are not fully understood or clearly defined, the resulting design may not meet the needs of the end users. This includes understanding the environment in which the asset will be used, as well as the desired performance, reliability, and safety requirements.
2. Designing in Isolation: If the design is done in isolation, without considering how the asset interacts with other systems or processes, it can result in a suboptimal design. This is particularly true for complex systems where the performance of the whole is more than just the sum of its parts.
3. Short-Term Cost Focus: Design decisions are often driven by initial cost considerations, without taking into account the total lifecycle cost of the asset. This can lead to designs that are cheaper to build or purchase, but more expensive to operate, maintain, or dispose of.
4. Lack of Experience or Expertise: Designing an asset requires a deep understanding of the underlying engineering principles, as well as experience and expertise in the specific type of asset being designed. If this is lacking, the resulting design may not be optimal.
5. Technological Limitations: Sometimes, the available technology may not allow for an optimal design. As technology evolves, what was once considered the best possible design may become suboptimal.
6. Regulatory or Standards Constraints: Sometimes, design decisions are constrained by regulatory requirements or standards that may not necessarily lead to the most optimal design.
7. Failure to Consider Maintenance: If the ease of maintenance is not considered during the design phase, it can lead to higher maintenance costs, more downtime, and a shorter asset lifespan.
8. Changes Over Time: Requirements, technology, and environmental conditions can change over time. A design that was optimal when it was created may become suboptimal as these changes occur.
9. Different entities; The need to subcontract

By addressing these issues during the design process, it's possible to create an asset design that is more optimal and delivers better value over its lifecycle.

Manufacturing and construction processes can be suboptimal due to several reasons:

1. Poor Planning: Inadequate planning can lead to inefficiencies, delays, cost overruns, and quality issues. Good planning involves a thorough understanding of the process, resources, time frames, and potential obstacles.
2. Lack of Skilled Labor: The lack of skilled labor can result in errors, inefficiencies, and lower quality output. Skilled workers can perform tasks more accurately and efficiently, reducing waste, improving quality, and decreasing the likelihood of rework.
3. Inadequate Quality Control: Without robust quality control processes, defects or errors might go unnoticed until they cause significant issues. Poor quality can result in rework, recalls, or failures during operation.
4. Outdated Technology or Equipment: Using outdated technology or equipment can limit productivity and efficiency, leading to higher costs and longer production times. It can also result in lower quality products or structures.
5. Inefficient Processes: Processes that are not streamlined or optimized can result in wasted time, materials, and effort. This includes unnecessary steps, inefficient layout of facilities, or not taking advantage of automation where appropriate.

6. Insufficient Training: Without adequate training, workers may not understand how to perform tasks correctly or use equipment effectively. This can result in mistakes, lower productivity, and safety issues.
7. Supply Chain Issues: Problems in the supply chain, such as delays in material delivery or poor quality of materials, can disrupt production schedules and affect the quality of the final product.
8. Lack of Safety Considerations: If safety isn't prioritized, accidents can occur leading to worker injury, legal issues, delays, and cost overruns.

By addressing these factors, manufacturing and construction processes can become more optimal, resulting in higher quality products or structures, lower costs, improved safety, and better overall efficiency.

Installation and commissioning are critical stages in an asset's life cycle. However, they can sometimes be inadequate due to various reasons:

1. Poor Planning: Insufficient planning can lead to oversights during the installation and commissioning phase. This can result in avoidable delays, cost overruns, or even incorrect installation.
2. Lack of Skilled Labor: If the team lacks skilled labor or proper training, they may make errors during installation and commissioning. This can lead to improper setup, which might cause operational issues or even equipment damage.
3. Insufficient Time: Sometimes, due to tight schedules or last-minute changes, the time allocated for installation and commissioning is not adequate. Rushed commissioning might overlook important steps or checks, leading to potential issues during operation.
4. Incomplete or Unclear Instructions: If the manufacturer's installation and commissioning instructions are incomplete, unclear, or not properly understood, this can lead to mistakes and potential issues later on.
5. Inadequate Testing: During the commissioning phase, the asset should be thoroughly tested to ensure it operates as expected under different conditions. If testing is inadequate, hidden issues may not be discovered until the asset is in operation, which can lead to failures or decreased performance.
6. Improper Handling or Storage: Mishandling or incorrect storage of equipment during the installation phase can cause damage, which might not be immediately apparent. This can lead to failures or malfunctions during operation.
7. Lack of Coordination: In larger projects involving multiple teams or contractors, lack of coordination can result in errors or oversights during the installation and commissioning phase.
8. Environmental Conditions: Sometimes, the site conditions where the asset is installed might not have been adequately considered. This could mean that the asset is not suitable for the environment in which it's placed, leading to early failure or decreased performance.

By ensuring adequate planning, skilled labor, time allocation, clear instructions, thorough testing, proper handling, coordination, and consideration of environmental conditions, the installation and commissioning process can be more effective, ensuring the asset performs as expected.

Maintenance processes can sometimes be suboptimal due to several reasons:

1. Reactive Maintenance Approach: Relying solely on a reactive maintenance approach, i.e., only fixing things when they break, often leads to unplanned downtime, higher repair costs, and potentially more damage.
2. Insufficient Preventive Maintenance: If preventive maintenance is not conducted regularly or effectively, it can lead to premature equipment failure. Skipping regular maintenance tasks can lead to minor issues becoming major problems.
3. Inadequate Training: Maintenance staff must be properly trained to carry out maintenance tasks and troubleshoot issues. Lack of adequate training can lead to mistakes, overlooked problems, or improper handling of equipment.

4. Poor Planning and Scheduling: Poorly planned maintenance activities can lead to increased downtime, inefficient use of resources, and missed maintenance tasks.
5. Lack of Proper Tools or Equipment: Without the right tools or equipment, maintenance tasks may not be performed effectively or safely. This can lead to inadequate maintenance and increased risk of equipment failure.
6. Not Following Manufacturer's Recommendations: If the maintenance team does not follow the manufacturer's recommendations for maintenance procedures and intervals, it can lead to premature failure or void warranties.
7. Ignoring Minor Issues: Small issues, if ignored, can escalate into major problems. An effective maintenance program identifies and addresses minor issues before they become significant.
8. Inadequate Maintenance Documentation: If maintenance activities are not accurately recorded, it can lead to repetitive errors, overlooked maintenance tasks, and lack of historical data for troubleshooting or predictive maintenance.
9. Lack of Predictive Maintenance: Predictive maintenance, which uses data to predict equipment failures before they happen, can reduce downtime and maintenance costs. If an organization isn't utilizing predictive maintenance where appropriate, they could be missing out on these benefits

Operation of an asset or a system can sometimes be suboptimal due to several reasons:

1. Improper Usage: If the equipment is not used as intended or according to manufacturer guidelines, it can result in decreased efficiency, increased wear and tear, and potential failures.
2. Insufficient Training: Operators who are not adequately trained may not use the equipment efficiently or may misuse it, leading to decreased performance or potential damage.
3. Ignoring Warning Signs: If operators ignore minor issues, unusual noises, or warning signals, small problems can escalate into more serious issues or even equipment failure.
4. Poor Environmental Conditions: If the equipment is operated under conditions it was not designed for, such as extreme temperatures, high humidity, or in an excessively dirty environment, it can lead to decreased performance or increased wear and tear.
5. Lack of Maintenance: Regular maintenance is crucial for optimal operation. If maintenance is neglected or not carried out properly, it can result in decreased performance, higher energy consumption, or premature failure.
6. Inadequate Resources: Lack of necessary resources, such as sufficient power supply, quality raw materials, or proper tools, can impact the operation of equipment or systems.
7. Operational Stress: Overloading or operating the equipment beyond its design capacity can lead to faster degradation, more frequent breakdowns, and shorter lifespan.
8. Inadequate Monitoring and Control: If the operation is not monitored effectively and adjustments aren't made as necessary, efficiency can drop and potential issues might go unnoticed.
9. Failure to Adapt: Changes in technology, regulations, or business requirements can make current operations suboptimal. Organizations need to continually adapt their operations to keep up with changes.

By addressing these issues, operations can be optimized, resulting in better performance, increased efficiency, reduced operational costs, and extended asset life.

Get Chat GTP to consider why things go wrong at each of these stages

7.9 RCM Misconceptions

Reliability Centered Maintenance (RCM) is a powerful methodology used to determine the most effective maintenance strategies, but it's often misunderstood. Here are a few common misconceptions:

1. RCM is Only for Large Companies: RCM is often perceived as a methodology only for large companies with vast resources. However, any organization, regardless of size, can benefit from the RCM approach if it is applied correctly.
2. RCM is Too Expensive: While RCM requires an initial investment in time and resources, it can deliver significant cost savings in the long run by increasing equipment reliability, reducing maintenance costs, and preventing expensive failures.
3. RCM Means More Maintenance: Many people believe RCM will always result in more maintenance tasks. However, the primary objective of RCM is to identify the most effective maintenance tasks, not necessarily to increase them. Often, it may result in fewer, but more targeted, maintenance activities.
4. RCM is Too Complex: Some people might think the RCM process is too complicated or time-consuming. While it's true that RCM involves a systematic and detailed analysis, the process is not beyond the grasp of maintenance teams who are given proper training and support.
5. RCM is Just About Maintenance: Although maintenance is a significant component of RCM, it also involves a broader understanding of how equipment fails and the consequences of these failures. RCM takes a holistic view of the system and is as much about understanding failure modes and improving design as it is about performing maintenance tasks.
6. RCM is a One-Time Process: RCM is not a "do it once and forget it" process. It's a continuous cycle that involves regular reviews and adjustments based on changes in operating conditions, equipment performance, and maintenance capabilities.
7. Every Asset Requires an RCM Analysis: While RCM is a powerful tool, not every asset requires a full RCM analysis. It is most effectively applied to critical equipment where failures can have significant consequences. For less critical assets, other maintenance strategies might be more cost-effective.
8. RCM is a type of maintenance
9. RCM is about CBM

By understanding these misconceptions, you can better apply RCM principles and reap the benefits of improved reliability and cost-effective maintenance

CHAPTER 8

The Benefits of RCM

8.1 The Early Benefits of RCM in the Aviation Industry

Reliability Centered Maintenance (RCM) has been widely adopted in the aviation industry since the 1970s. The implementation of RCM in the aviation industry has brought significant benefits, including improved safety, reduced maintenance costs, and increased equipment reliability.

One early example of the benefits of RCM in the aviation industry is its implementation by United Airlines in the 1980s. United Airlines used RCM to identify the most critical equipment on its fleet of Boeing 747 aircraft and develop appropriate maintenance strategies. As a result, United Airlines was able to reduce its maintenance costs by over $200 million per year while still maintaining a high level of safety and reliability.

Another example is the implementation of RCM by the US Air Force in the 1990s. The US Air Force used RCM to identify critical equipment on its aircraft and develop appropriate maintenance strategies. As a result, the US Air Force was able to reduce the number of aircraft mishaps by 80% while reducing maintenance costs by over $4 billion per year.

The table below shows flight accident rates per year and per 1000000 km, excluding terror attacks:

Year	Total Flights	Total Distance (million km)	Fatal Accidents	Fatalities	Accident Rate (per 100,000 flights)	Accident Rate (per 1,000,000 km)
1960	1,620,000	4,200	47	1,535	2.90	11.19
1961	1,780,000	4,800	45	1,356	2.53	9.38
1962	1,980,000	5,500	49	1,493	2.47	8.91
1963	2,230,000	6,300	42	1,328	1.88	6.80
1964	2,470,000	7,400	46	1,324	1.86	6.22
1965	2,740,000	8,700	50	1,262	1.82	5.75
1966	3,050,000	10,300	49	1,525	1.61	4.75
1967	3,370,000	12,400	55	1,323	1.63	4.43
1968	3,770,000	14,700	51	1,202	1.35	3.46
1969	4,150,000	17,200	52	1,191	1.25	3.03
1970	4,670,000	20,300	46	1,063	0.98	2.26
1971	5,160,000	23,100	41	1,096	0.80	1.78

1972	5,770,000	27,100	48	2,374	0.83	1.77
1973	6,420,000	30,500	50	1,863	0.78	1.64
1974	6,950,000	34,300	49	1,261	0.71	1.43

Year	Total Flights	Total Distance (million km)	Fatal Accidents	Fatalities	Accident Rate (per 100,000 flights)	Accident Rate (per 1,000,000 km)
2015	34,986,270	91,423.8	16	560	0.046	0.175
2016	35,476,088	94,773.5	19	325	0.054	0.201
2017	37,795,360	101,313.2	10	44	0.026	0.099
2018	38,795,548	107,036.4	11	523	0.028	0.103
2019	39,891,000	113,403.6	8	297	0.020	0.071
2020	23,921,253	61,411.7	5	299	0.021	0.081

Source: Aviation Safety Network (https://aviation-safety.net/)

8.2 Benefits in other industries

TBC

8.3 Learning/ Understanding

The "known knowns" model, popularized by former U.S. Secretary of Defense Donald Rumsfeld, is a framework that categorizes elements of knowledge and uncertainty into four distinct quadrants:

1. Known Knowns: These are things we know that we know. They represent familiar, understood information within our expertise or experience.
2. Known Unknowns: These are things we know that we don't know. They acknowledge gaps in our knowledge that can potentially be filled through further research or inquiry.

3. Unknown Knowns: These are things we don't know that we know. This category often includes subconscious or tacit knowledge, which we may not realize we possess until it's called upon.
4. Unknown Unknowns: These are things we don't know that we don't know. They represent unforeseen challenges or completely new information that couldn't have been anticipated based on existing knowledge.

The best quadrant to be in is the known knowns. The worst is the unknown unknowns.

This is very relevant to an RCM analysis. When RCM is applied properly, a group of people who each know a bit about the equipment answer relevant and searching questions about how the asset works, how it can go wrong, what happens if it goes wrong and what can be done about it.

The group may often be asked questions they don't know the answer to. This is good news because the group is finding out what they do not know. In other words they are identifying the unknown unknowns. These are the most dangerous. But simply by identifying what they do not know is progress and the unknown unknown has become a known unknown. The next step, of course, is to find the answer, and move to the most favourable quadrant.

None of the questions are superfluous or unnecessary – they all need to be answered if we are going to manage our assets efficiently and effectively.

It is uncanny how often – when RCM is done properly – people with great experience of the equipment, come away and say they have learnt a whole lot more about it.

One recent example from the author's experience was a Pharmaceutical company. The fifth version of a specialist machine, used to manufacture a well-established drug was being analysed. The machine was at the design stage and the Pharmaceutical company versy sensibly carried out RCM analysis on the machine with their best operators and maintenance technicians, and paid the for time of the machine designer from another company. A large number of risk mitigations were identified and incorporated into the new design CJ blue box example

In this way, RCM is used as a learning tool, and a very effective one at that. Improved learning and understanding is a 'soft' benefit, but this rapidly translates into 'hard' benefits such as improved cost, availability, product quality, safety and reliability etc.

Whoever asks the operators (or the maintenance technicians for that matter)? The above example highlights another valuable aspect to RCM. It asks the people at the coal face. Like all professions, there are good operators and maintenance technicians, and some are not so good.

But the knowledge and insights from the good ones are often worth their weight in gold. The problem is – in many cases - nobody ever asks them.

One famous industrialist known for his practice of speaking to people on the shop floor was Henry Ford, founder of the Ford Motor Company. Ford was known for his hands-on approach and often visited the production lines of his factories. He believed in the importance of understanding the needs and concerns of his employees, and used these insights to make decisions that would improve both working conditions and manufacturing efficiency. He was instrumental in introducing the concept of the assembly line in automotive production, which revolutionized industrial manufacturing.

RCM standards

JA1011, titled "Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes," is a standard published by the Society of Automotive Engineers (SAE). This standard provides the minimum criteria that a process must possess to be RCM.

The main focus of JA1011 is to outline the basic requirements for an RCM process, specifically dealing with the structure, analysis tasks, and steps that should be present. It includes seven questions that any process must answer to be considered RCM:

1. What are the functions and associated desired standards of performance of the asset in its present operating context?
2. In what ways does it fail to fulfill its functions?
3. What causes each functional failure?
4. What happens when each failure occurs?
5. In what way does each failure matter?
6. What systematic task can be performed proactively to prevent or to diminish to a satisfactory degree the consequences of the failure?
7. What must be done if a suitable preventive task cannot be found?

By addressing these questions, an organization can apply the RCM methodology effectively to improve the reliability and maintenance practices of their assets. The standard does not specify how to perform these tasks, which allows flexibility for different industries and types of equipment.

SAE JA1012, titled "A Guide to the Reliability-Centered Maintenance (RCM) Standard," is a document published by the Society of Automotive Engineers (SAE). This guide provides instructions for applying the processes described in the SAE JA1011 standard "Evaluation Criteria for RCM Processes".

Whereas JA1011 outlines the minimum criteria a process must meet to be considered as conforming to the principles of RCM, JA1012 provides detailed guidelines and methods to apply those principles. The standard is designed to be usable by any

> industry and for any product, whether it be an aircraft, a manufacturing facility, a railway, a vehicle, etc.
>
> SAE JA1012 assists users in understanding the seven questions outlined in JA1011, helping them in defining failure modes, identifying and analyzing failures, deciding on applicable maintenance tasks, and more. It is a valuable resource for anyone implementing a reliability-centered maintenance program in their organization.
>
> IEC 60300-3-11:2009 is a standard issued by the International Electrotechnical Commission (IEC). It is part of the IEC 60300 series, which is specifically focused on dependability management. The "3-11" portion represents a specific part of this series.
>
> IEC 60300-3-11:2009, titled "Dependability management - Part 3-11: Application guide - Reliability Centred Maintenance," provides guidelines for implementing Reliability Centred Maintenance (RCM) to optimize maintenance strategies.
>
> The standard provides guidance on:
>
> 1. How to apply RCM in the development and implementation of maintenance plans.
> 2. How to integrate RCM with other dependability activities.
> 3. The necessary activities for the successful implementation of RCM.
>
> The principles of RCM that are outlined in this standard include the understanding and management of the changing failure modes associated with assets over their lifecycle, identifying appropriate maintenance strategies for those assets, and making decisions about maintenance plans.
>
> The standard is applicable across industries, and can be used by organizations seeking to optimize their maintenance strategies for dependability, cost-effectiveness, and a focus on critical assets.
>
> Please note that standards are updated periodically and it is important to refer to the most recent version or amendment. For the most accurate and current information, please refer to the official IEC website or other authorized sources

These standards are consistent with each other.

Glossary

Analysis Selection Process

An algorithm containing criteria unique to an organization used to determine which analysis process to apply to an asset: RCM, REM, LCS, or RTF. RCM, REM, LCS, or RTF.

Conditional Probability of Failure

The probability that a failure will occur at a specific age once the item has survived to that age.

Default Strategy

A failure management strategy, other than proactive maintenance, implemented to manage the con sequences of failure (e.g. procedural check, Failure Finding task, no scheduled maintenance, engineering redesign).

Design Capability

What an asset, or any part thereof, is capable of doing. Essence of RCM: To manage the consequences of failure.

Evident Function

A Function that, upon its loss, becomes evident to the operating crew under normal conditions.

Facilitator

An individual who is very well versed in RCM and principles, leads the working group and validation team, ensures that RCM principles are being applied correctly.

Failure caused by a Failure ConsequencesMode

mattersDescribe; the how the categories loss of of function Failure Consequences are safety, environmental, operational, and non-operational.

Failure Mode

What specifically causes a Functional Failure. Failure that records Modes the first and four Effects steps of Analysis the (RCM FMEAprocess): A:document FunctionsEffects., Functional Failures, Failure Modes, and Failure Failure Modes, Effects, and Criticality zilkAnalysis (RCM Failure FMECAprocess): EffectsA:, and document FunctionsFailure, that Functional records Consequencesthe Failures.first, five Failure steps of Modesthe, Failure that records Modes the first and four Effects steps of Analysis the (RCM FMEAprocess): A:document FunctionsEffects., Functional Failures, Failure Modes, and Failure Failure Modes, Effects, and Criticality zilkAnalysis (RCM Failure FMECAprocess): EffectsA:, and document FunctionsFailure, that Functional records Consequencesthe Failures.first, five Failure steps of Modesthe,

Failure nothing were Effect

done A to brief predict or description prevent of what the would Failure happen Mode.if Failure vice is in a Finding failed state Task(: e.g. A task blowing that smoke checks at if a a smoke protective de detector to ensure that it sounds an alarm).

Failure the consequences Management of a Failure StrategyMode

An or a action Multiple taken to Failure (manage e.g. proactive redesign, operating maintenance, procedureFailure).Finding task, physical

Final that reached details Summarybetween the

final the A report results working of that the group is RCM and compiled the analysisonce.validation consensus team is

Function

Describes the performance required of an asset.

Functional Failure

An unsatisfactory condition in which either some or all of a Function cannot be performed.

Hidden Function

A Function that, upon its loss, does not evident to the operating crew under normal become conditions.

Information Worksheet

A document in which Functions, Functional Failures, Failure Modes, and Failure Effects are recorded.

Lubrication, Cleaning, and Servicing (LCS)

An analysis process limited to formulating lubrication, cleaning, servicing, and Failure Finding tasks that may be applied to a physical asset if the rigor of RCM and REM are not warranted. Multiple Failure: Includes the failure of a protective device and another failure (e.g., low pressure switch fails and system pressure falls below normal levels).

Net P-F Interval

The minimum time remaining to take action in order to manage the consequences of failure once an On-Condition task is performed.

On-Condition Task

A proactive task performed at a defined interval to detect a potential failure condition so that maintenance can be performed before the failure occurs.

Operating Context

A document that includes a storybook identification of the system to be analyzed.

P-F Interval

The time from when a potential failure condition is detectable to the point that failure occurs.

Partial Failure

The inability to function at the level of performance specified as satisfactory.

Potential Failure Condition

Evidence that a Failure Mode is in the process of occurring (e.g., vibration, heat, cracks).

Preventive Maintenance

A scheduled Restoration or scheduled Replacement task that is performed at a defined interval without considering the item's condition at the time of the task.

Primary Function

The main purpose an item or system ex ists.

Proactive Maintenance

Includes scheduled Restoration, scheduled Replacement, and On-Condition tasks.

Procedural Check

A task performed at a specified interval to check for an evident failure that may have already occurred because it would be unacceptable to operate or continue operating with the failure condition. Protective Device: A device or system intended to protect people, the asset, and the organization in the event that another failure occurs (e.g., smoke detector).

RCM Failure Decision Modes as Diagramevident

or An hidden, algorithm assess used Failureto classify ConsequencesDefault technically Strategy, appropriate is determine and recommendedif any worth.proactive doing, and maintenance determine is if a ism Reliability strategies structured process required Centered to used to ensure identify an Maintenance asset the (meets RCMfailure)its: A mission zeromanagement -based, requirements and cost-effective in its manneroperation-al.environment in the most safe Reverse Engineering Method (REM): An analysis process limited to analyzing protective devices and reverse-engineer ing manufacturer recommend-ations or a suitable template (such as a legacy maintenance program) to identify suitable failure management strategies.

Scheduled Replacement Task

A task performed at a spec i6ed interval that replaces an item without considering the item's condition at the time of the task.

Scheduled Restoration Task

A task performed at a speci fled interval that reworks or restores an item's failure resistance to an acceptable level without considering the item's condition at the time of the task. Secondary Function: A function, other than the primary function, of a system.

Total Failure

Complete loss of Function.

True RCM

An RCM process that embodies the seven steps of RCM, which are carried out in order, as defined by SAE JA1011, Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes.

Useful Life

The age at which a significant increase in the conditional probability of failure occurs.

Validation Meeting

A meeting organized to review the working group's decisions to ensure they are safe and technically defensible and to determine what recommendations will be implemented.

Validation Package

A report prepared by the facilitator that contains all documentation necessary to perform a technical review of the working group's recommendations.

Working Group

A team of equipment experts assembled to carry out the RCM analysis; team members represent various disciplines such as a mechanic, operator, original equipment manufacturer (OEM), and a systems engineer.

Zero-Based RCM Process

The RCM process is carried out assuming that nothing is being done to predict or prevent Failure Modes.

References

Source list as given in the original handout.

- Mark book
- Moubray
- Smith DJ
- JA1011
- JA1012
- IEC
- OREDA
- Nancy book
- Gehris book