RCM Changed My Life

This blog documents my RCM experiences, starting with my first introduction to RCM in 1994

We are now at the beginning of day 2 of my 3-day RCM course, moving on to the 4th of the 7 RCM questions which asks about the effects of each failure mode.

We learned that it is important to write down what would happen when the failure occurred if nothing was done to predict, prevent or detect the failure. That seemed odd at first but I was quickly comfortable with it after a little thought. By writing what would happen if there was no maintenance to manage the failure, you actually get a full statement of the problem. If we don't do any maintenance (or we do the wrong maintenance), then these undesirable things will happen. So, we have now written down a large proportion of the justification for any maintenance tasks (or, indeed, the justification for the doing no maintenance). If we do decide on a maintenance task then we know exactly why because we can see which undesirable events it is going to prevent. That just made so much sense to me.

We also learned what should actually be described in the failure effects:

- what evidence there is (if any) that the failure has occurred
- in what way it affects safety and/or the environment (if it does at all)
- how it affects production or operations
- any secondary damage
- what must be done to repair the failure.

So, you end up with a short story that succintly describes everything between the point at which the failure occurs and the point at which the equipment is back into service. That in itself was very refreshing to me. As a reliability engineer I was used to seeing FMEAs that contained frustratingly skimpy effects like "pump stops", "loss of sensor output" or "invalid temperature signal" which left the reader to work out for themselves what else happens.

Correctly written, RCM failure effects provide the reader with a full account of what happens. This helps to rapidly provide a complete understanding of the equipment and how it interacts with its environment. I have since seen how this additional knowledge boosts the confidence of technicians and engineers, especially when it comes to fault finding. One fitter still stands out in my mind after many years. He participated in the analysis of an HVAC system and was quite concerned that his background (mainly as a diesel fitter with no air conditioning experience) would mean he would be able to contribute little to the analysis and that his presence would prove to be a waste of time. He was actually there because the system included a diesel powered heater and he revelled in explaining it to the rest of the group - and the heater failure effects were consequently excellent. By the time we had finished the failure modes and effects for the air conditioning side of the system, this fitter suddenly understood the principles of air conditioning and he felt confident to get fully involved with it in the future (he previously would have let someone else deal with it).

In my opinion, a proper explanation of what happens to a piece of equipment (and the system it belongs to) when a failure occurs is "worth its weight in gold". (more)
Posted by plumothy on Aug 25, 2009 9:52 AM GMT

RCM Question Three

We are now approaching the end of day one of my first RCM course and it is time to learn about failure modes, which are the subject of the 3rd of the 7 RCM questions: What causes each functional failure?

I started to feel comfortable at this point. I knew what failure modes were - I had already spent many years maintaining against them, diagnosing them, repairing them, and analysing them. I already knew that when a failure mode occurred on a piece of equipment, it put it into a state where it was no longer capable of performing its intended function(s).

But I had the classic engineering blinkers on. As a reliability engineer I thought only of the equipment and what could go wrong with it. If you spend long enough looking at component-based FMEAs then I guess your brain gets conditioned to "failure mode equals broken component".

So, I enjoyed learning that RCM regards a failure mode as any event which can cause a functional failure. This, therefore, includes human error, the design of the equipment and the way in which it is being used. It is not limited to broken components.

My reliability engineering brain thought of another concern. I was used to seeing FMEA with hundreds or even thousands of failure modes listed at a minute level of detail (down to individual components on electronic boards for example). Wouldn't it take for ever to analyse all of these?

Well, I needn't have worried. RCM asks you to list failure modes in no more detail than is necessary for you to later be able to use the decision logic to determine how to manage them. So, if a card has 50 components on it and you never investigate which one of them causes the card to fail (you just replace the entire card) then you can say "card fails".

In addition, RCM asks you to list only reasonably likely failure modes. Whilst "meteorite strikes pump" is a possibility, it is such a slim possibility that you would not write it down.

We also saw that think of failure modes as falling into the following 3 categories was an aid to ensuring that no important ones would be missed. You should list:

- failure modes that have occurred before
- failure modes that existing maintenance is preventing
- failure modes which don't fit the 1st 2 categories, but you still think they are a real possibility.

Finally, I liked the way failures could be grouped according to the reasons why they occur:

- gradual deterioration of the equipment (and therefore its capability) due to normal wear and tear - e.g. the moving parts of an engine wear out
- overloading the equipment (expecting more of it than it is capable of) - e.g. you drive your car too fast round a corner
- the equipment was actually built and put into service when already in a failed state - e.g. the Hubble Space Telescope mirror. (I won't mention which one it was, but I knew of an entire train that fell into this category!)

  (more)

Posted by plumothy on Jul 24, 2009 3:15 PM GMT
RCM Question Two

It's now after lunch on day one of my RCM course. We've seen the importance of writing a good set of functions with proper performance standards. Now it's time to move onto RCM Question 2: Functional Failures.

Question 2 actually asks **how** the equipment could fail to perform its functions, not **why**. Once this important distinction was understood, writing functional failures became trivial. It is simply a matter of writing down the opposite of the entire function or part of it (total failure or partial failure).

Simon Deakin, our tutor, pointed out that it was such a trivial step that we might wonder why it was included in the RCM process at all. In my opinion there are two very good reasons:

1. Each functional failure is a description of one of the equipment's failed states. Having written functional failures, everyone now knows what is actually meant by failed - and that's not as obvious as you might think. Containment of oil by a gearbox seal was a good example used on the course. When has the gearbox failed to perform the function "to contain the oil"? Is it when the first drop gets past the seal? Is it when there is a hazardous pool of oil on the floor? Is it when the cost of topping up the oil becomes too high? Or is it when so much oil has leaked out that the gearbox seizes up? These are 4 different functional failures (each actually expressing a different performance standard - which was missing from the function). It doesn't matter which one of them the RCM review group chooses, but once they have chosen it then we can determine if we should a) check frequently for the tiniest leak, b) check for slight leaks and inspect the floor, c) monitor oil consumption or d) just keep pouring in the oil so it doesn't seize.

2. The very first question on the RCM2 Decision Diagram (Question H) starts "***Will the loss of function caused by...***". The RCM review group is being asked about the loss of function (i.e. the functional failure) not the failure mode itself. I believe this question is answered incorrectly more often than any other on the decision diagram. Clearly a big part of that is because it is the only question which has to be answered for every failure mode. But, in addition to that, it can often be a difficult question to answer, sometimes requiring the group to think very carefully about the visibility of the failure mode versus the evidence of the functional failure (it is easy to confuse the two). So, a well-written functional failure (coupled with careful wording of the failure effects) makes the RCM facilitator's job much easier when he/she is asking for the answer to Question H.

So, that was about it for function failures. They are really quick and easy to write - then you move on.
(more)

Posted by plumothy on Jul 22, 2009 8:33 AM GMT

RCM Question One

It’s still day one of the course. Having looked at the background and history of RCM, we now moved on to the first of the seven basic questions.

We has already seen that maintenance means keeping the equipment doing what its users want it to do. So that’s why RCM starts off by asking what the users of the equipment want it to do. It sounds like an easy question to answer, but I realised there’s more to it than you might think.

Ask a non-RCM person what they want their car to do and they will probably say something like: “to get from A to B”. But your feet can do that.

So, they might amend it to: “to get from A to B at speeds up to 70mph”. But a helicopter could do that.

Ok, how about: “to get from A to B at up to 70mph along made up roads”. But a motorcycle can do that - what about passengers and luggage?

The next iteration might be: “to carry up to 4 people from A to B at up to 70mph along made up roads”.

That trivial example illustrated to us that you need to think not only about what the equipment should do, but also about how well it should do it.

Experienced RCM facilitators think like this automatically. But when you are doing it for the first time, it can be tricky. It is often difficult to get your engineering brain out of the mode in which it wants to describe what the equipment should **be** and into the mode where it can describe what the equipment should **do**.

I had already seen examples of this on the railways. New train specifications often contained a lot of statements about what the train should be, mixed with some about what it should do. One clause might say the train should be capable of accelerating from 0 to 90mph in 3.5 miles. And the next clause might say it should have GEC traction motors.

I wondered what a train specification might look like if it contained only functions, ie just a list of things the train should be able to do. Years later, I got a feel for this when I was watching a TV documentary about 9/11. The Pentagon was being refurbished at the time the plane crashed into it. A Pentagon official told us that the specification for the refurbishment was effectively just a list of functions, ie a statement of what the building should do. The previous refurbishment many years earlier had a more traditional specification and described in enourmous detail what the building should be (down to where every electric socket and light switch should be positioned, for example). The old style of specification occupied **several thousands pages**. The new functional specification was **less than 20 pages**.

My first attempt at writing RCM functions was in an exercise on the course. We had to write functions for the 35mm slide projector being used for the course. I enjoyed forcing my brain to think in a different way because I was looking forward to where it would take us next. I have just had a look at that function list, expecting to cringe at what I had written. However, they're not too bad considering it was my first time.

So, I'm a fan of functions (although I know some people think writing them is a waste of time). It doesn't take long to write them. They are the foundation for the rest of the RCM analysis and give the RCM review group a more complete sense of what the equipment is about. A real analysis that illustrated this very emphatically for me a few years later was a beer filter in a brewery. We very quickly wrote function 1 which was about removing particle matter from the beer to a particular standard of clarity. Then I asked what else the filter does. Silence! They thought that was it. So, I probed deeper and, only half an hour later, we had 17 functions written down. Without those functions I know we would have missed some failure modes, two of which proved to be extremely important.
(more)
Posted by plumothy on Jul 17, 2009 7:21 AM GMT

Same Component, Different Use, So Different Maintenance

The next thing that really caught my attention in my RCM education was being shown how the **same piece of equipment** could require a very **different set of maintenance tasks** depending how it is used. Or, in other words, it depends on the equipment's **operating context**.

The term "operating context" is used at the end of the definition of RCM:
A process used to determine the maintenance requirements of any physical asset in its operating context.

So, before you begin to answer the 7 basic RCM questions, you should make sure that you understand the equipment's operating context.

On my RCM course, an example using 3 pumps illustrated very simply how operating context can greatly influence the maintenance that needs to be done.

The 3 pumps were absolutely identical, ordered from the same manufacturer at the same time - and they all came with the same manufacturer's maintenance manual. The manual said that the user should perform vibration analysis on the impeller shaft bearings each year, inspect seals for leakage every month, overhaul the impeller every 5 years and so on.

This seemed like a a reasonable set of maintenance tasks for **Pump A**. Pump A had to pump cooling water from Tank X to Tank Y at 1,000 litres per minute. If it fails then the downstream process will have to stop when Tank Y runs dry at a cost of thousands per hour.

There is also **Pump B**. It is on a similar production line to Pump A and also must pump 1,000 litres per minute. The difference is that Pump B is a duty pump with **Pump C** available next to it as a standby. So, it does not matter if Pump B fails because Pump C can take over (no lost production). Hence, it may be difficult to justify doing any maintenance on Pump B because it does not cost any lost production when it fails. So, for a lot of its failure modes (and possibly all of them) fix on failure might be the right thing to do.

Ah, but what about Pump C? If it fails (while Pump B is still running) then it doesn't matter (the production process still continues). But, because we need to know that Pump C will work when Pump B fails, we could test Pump C at regular intervals to see if it is still working (and fix it if it isn't). In RCM this is known as "failure finding".

So, **these three identical pumps** could end up with three very **different** sets of **maintenance** tasks as a result of their individual operating contexts.

I found this to be very illuminating and quite fascinating. The manufacturers can only write a generic set of maintenance tasks to suit how they envisage their equipment might be used - a lot of the time they are going to be wrong.

Later on in my RCM career I enjoyed a spectacular example of this when I was facilitating an analysis for a client. The equipment supplier was in the review group jealously guarding his 500 failure mode FMECA. During the analysis we were discussing a failure that caused a pressure switch to trip in a protection system. The supplier "blew a fuse" when it dawned on him that two thirds of his FMECA was completely wrong (along with most of the reliability predictions based upon it). They had assumed that their customer would wire the pressure switch up to trip on falling pressure. But, they actually wired it up to trip on rising pressure (and for serious safety reasons this could not be changed). It was the equivalent of learning that standby Pump C does not actually exist in the example above - but 10 times worse!

After he calmed down, the supplier realised that his presence in the RCM review group had just saved his company serious problems and embarrassment in the future.

So, it was an important lesson that I enjoyed learning: you need to understand the equipment's operating context before you can begin to determine the correct maintenance tasks for it.

Posted by plumothy on Jul 14, 2009 11:56 AM GMT

The Seven Questions

☆☆☆☆☆

Still only part way through my 3-day RCM course. I now know that on the railways (as, presumably, in most other industries also) we have probably been causing most of the failures through our own ignorance.

Having been knocked for six with that bombshell, I was looking forward to learning what RCM suggests we actually do about it to put things right.

The answer came in the logical sequence of steps known as the **seven basic questions**. Answer these questions honestly and accurately and you will end up with as near as you can get to the perfect maintenance schedule for your equipment and the way you are using it. That last bit is important. The environment in which you use your equipment and the way in which you use it are known as its operating context. Put identical pieces of equipment into different operating contexts and they are likely to require different maintenance. In fact, operating context is so important that it is in the definition of RCM:
*Reliability-Centred Maintenance: a process used to determine the maintenance requirements of any physical asset in its **operating context**.*
RCM is applied by answering the following questions:

1. what are the functions and associated performance standards of the asset in its present operating context?
2. in what ways does it fail to fulfil 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 can be done to predict or prevent each failure?
7. what should be done if we cannot predict or prevent a failure?

These questions are the structured framework that form the basis of RCM. I will explain how I learned about each question in future posts.
(more)
Posted by plumothy on Jul 8, 2009 10:05 AM GMT
Why Do We Do Maintenance?

So, there I am, still in the middle of my introduction to RCM. Having seen the explanation for why maintenance can actually cause the failures it is meant to prevent, I was now in for another surprise.

The next thing this RCM course taught us was that equipment maintenance is not done **just to prevent failures**. Hmm, that didn't seem right at first; I had spent my career on the railways up to that point trying to prevent train failures; surely, that's why we were doing the maintenance, wasn't it?

Yes and no.

RCM's viewpoint is that we only do maintenance in order to keep the equipment doing **what the user wants it to do** - in other words we are making sure that the equipment can keep performing its **intended functions**.

You don't have to think too hard for too long to realise that this is actually true. A railway example comes to mind. Trains have solid steel wheels which rotate along solid steel rails. Both the wheels and the track receive regular inspections - scheduled maintenance in other words. This maintenance is not done in order to prevent the wheels and the track from failing; In fact, it can't do that because the they both inevitably wear each other out (especially on corners). The maintenance does not prevent that wear and tear - the failure mechanism continues regardless of the inspections. The maintenance simply monitors the progression of the failure and tells us when we should take some remedial action so that the failure is not allowed to reach its full conclusion (which would involve a derailment and even fatalities).

Train wheels are reprofiled on a lathe so that they can continue to guide the train along the rails. But, that can only be done so many times before the wheel has to be scrapped and replaced with a new one. Similarly, the rails have to be replaced before they completely wear out.

So, all that maintenance and remedial action **does not prevent the wheel or the rail from failing** - they both eventually end up in the **scrap yard**.

But, all that maintenance did keep the train running along the track; **IT KEPT THEM DOING WHAT THE USERS WANT THEM TO DO**.

This is why the RCM process starts off by asking: what are the functions of the equipment (or what do the users want it to do)?

The list of functions becomes a statement of the objectives of the maintenance schedule. In other words, having written a list of functions for a piece of equipment, we now know why we need to maintain it.
(more)
Posted by plumothy on Jul 6, 2009 9:55 AM GMT
I Was Shocked

It was 26th January 1994 - day 1 of my 3-day RCM course, run by Simon Deakin of Mutual Consultants Ltd.

We were being told how it all started in the civil aviation industry. In 1960 they were experiencing **60 crashes per million take-offs**, and **40** of those were due to **equipment failure** (not pilot error).

I quickly tried to do some mental arithmetic - just at Heathrow, if you assume 1 take-off every 3 minutes, that's about **1 crash per month**.

Simon explained that the airlines realised this was not good enough, so they decided they should do their scheduled maintenance (like engine overhauls) more often. **The crash rate increased**.

I was fairly surprised to hear that. More maintenance meant more failures which meant more crashes.

But, I was quite shocked when I heard the next bit. In 1960 we all thought that all equipment behaved according to what we now call Failure Pattern B: as equipment gets older, it becomes more likely to fail. That seemed quite reasonable to me - as it did to the airlines. They were maintaining all their equipment based on that assumption.

However, having got more failures when they did more maintenance they were puzzled, so they did an enormous amount of research into their failures. This research showed them that there were actually 6 failure patterns:

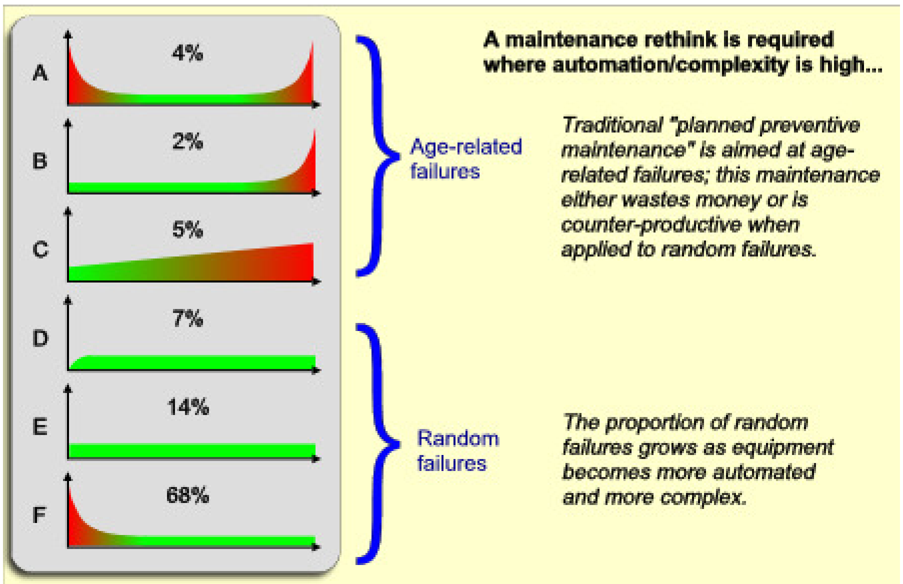

Worryingly, the research also showed them that **only 2%** of their equipment actually did fit Failure Pattern B.

And the real shock to them was that **68%** fitted Failure Pattern F (where disturbing the equipment while maintaining it actually **introduces failures**). It was also a big shock to me.

But why was it such a shock? It slowly dawned on me that I (and a lot of other railway colleagues) had experienced this before. We knew that failures on a train were more likely just after it had a scheduled exam. We knew that we couldn't allow a refurbished train straight into service without having a good look at it first ("post-works exam"). But we didn't know why.

At last, here was something logical which explained why this was happening on the railways (and, presumably, everywhere else). Even after a couple of hours, this RCM course was making a lot of sense. I couldn't wait for the rest of it to learn what RCM could do to help us stop causing so many Pattern F failures.
(more)

Posted by plumothy on Jul 3, 2009 9:02 AM GMT
Where It All Started

About 18 months into my time as a reliability engineer with British Rail I was sent on a training course that sounded like it might be quite interesting - 3 days of Reliability-Centred Maintenance (or RCM).

I had no idea what it was and it sounded like it would require tons of historical failure/reliability data. That was the only thing that I thought would put me off RCM - because I was beginning to get tired of manipulating masses of train failure data.

I needn't have worried! I was hooked before lunch on day 1. Suddenly, I knew why I had never been fully comfortable with:

- the way we maintained trains
- collecting and analysing records of thousands of failures that kept happening again and again
- our dismal attempts at improving train reliability.

For the rest of that course the hairs stood up on the back of my neck as I thought "wow, we've being doing it wrong for the last 150 years". This was going to make a real difference. I couldn't wait to get stuck into applying RCM and getting the railways working reliably.
(more)
Posted by plumothy on Jul 2, 2009 3:00 PM GMT