
Trouble in Store

RCM Notes series

Dr Mark Horton, Numeratis.com, September 2013

1 Introduction

I want to start by asking three questions.

- 1 When was the last time you saw someone draw a spare part from a store, install it, and then discover that it didn't work or it failed almost immediately?
- 2 When was the last time that your organisation reviewed the maintenance of key equipment kept in store?
- 3 Do you think that the answers to the first two questions could be related in some way?

2 Out of sight...

It is not surprising that engineers are keen to use tools like RCM to manage critical equipment. RCM is formalised common sense: it identifies failure modes and makes sure they are properly managed, either through maintenance or by making one-time changes to operations or design. The outcome is a set of maintenance schedules that reduce downtime, minimise maintenance and deliver the best possible performance. It's all career-enhancing work.

Perhaps it's because we think of critical equipment *doing* something—pumping, boiling, conveying, monitoring, containing, protecting and more—that our vision seems to cloud over when we think about the store that holds equipment and spare parts that aren't in use. But the truth is that if low demand equipment and parts aren't looked after in store, they won't do their job when they're needed. Our expensive, operations-based RCM analysis is just a heap of used paper if parts aren't ready to go when they're deployed.

3 Trouble in store

If it is so important to apply the right maintenance to parts and equipment in stores, can't it just be covered in the RCM analysis?

Sometimes it can. If the part is simple enough—perhaps it has half a dozen failure modes that could develop in storage—then they could be tackled as part of the operational RCM analysis. The problem with this approach is that “in store” is a completely different operating context from the operational scenario. The RCM review group's minds are focused on working equipment, and making the switch to a totally different context can disrupt the analysis process.

“The consequences of a failure in the storeroom are completely different”

The store's different operating context has a number of important consequences.

- 1 Most of the failure modes listed in the operational analysis won't occur in the warehouse. These include failures that result from mechanical motion, stress, wear, erosion, high temperatures and pressures and a host of other conditions that the part only experiences when it is in use.
- 2 A whole host of other failure modes will *only* happen in the store: slow deterioration of fluids, brinelling of bearings, structural distortion, damage from mishandling, seizing up of actuators and a host of others.
- 3 The *consequences* of a failure in the storeroom are probably completely different from those of a failure in use.


The third point is central to the way failures in store have to be managed. The RCM consequence

categories that are assigned to operational failures and failures in the store are not the same. For example, failure of a duty generator winding is almost certainly going to be classified as *operational*. The same failure mode in a storeroom would be effectively *hidden* until the generator is put into use. With a few exceptions such as leaking lubricants and some accidental damage, almost all storeroom failures will be completely hidden until the part is needed.

The difference in consequence categories, hidden versus operational, tells us something immediately: it's very likely that the maintenance requirements in the two operating contexts will be totally different.

4 Hidden failures bite (eventually)

The graph below illustrates failure rates in an ideal world. If its maintenance is properly carried out, a part will probably experience a low rate of failures during its operating life and none before it is installed.




As a concrete example, suppose that a pump set has an operational mean time between failures of five years. The chance of a failure every day is tiny:


$$1/(365 \times 5) = 0.00055 = 0.055\%$$

So most days are comfortable: during the equipment's lifetime, there is a small chance—around one twentieth of one percent—that a failure will occur on any day.

Of course the implicit assumption here is that the item does not fail while it is in the store.



But as we have already seen, the idealistic picture is wrong: equipment in stores fails; it just fails differently. The failure modes may be not be the same, the failure rate could be lower or higher, and the consequences will almost certainly be different, but if no maintenance or testing is carried out, the part will eventually fail.



Suppose that all the failures in store contribute to an overall pump set mean time between failures of 10 years; in other words, the failure rate in stores is half that of an operational pump set:

$$1/(365 \times 10) = 0.027\% \text{ per day}$$

Day-to-day, the failure rate is even less worrying than it is during operation. But most failure modes in the store are *hidden* until the part is needed. Although 0.03% per day seems insignificant, no one notices if the pump fails, and no one knows that it needs to be repaired. What happens is that all these daily probabilities of failure in store get added together until the part is finally used.


“The real-life experience is... a very uncomfortable pattern of early life failures”

Suppose that the pump set is held in store for six months before it is used. The total chance of

failure on the day it is withdrawn from the warehouse and installed is about¹

$$0.5/10 = 5\%$$

So the chance that the newly-installed pump fails immediately is about 100 times the normal, daily failure rate.



That is why installing a replacement part can be a white-knuckle experience: hidden failures accumulated over the whole period in store suddenly become evident when it is called on to operate. The real-life experience isn't a comfortable low rate of random failures (Nowlan and Heap pattern E), but a very uncomfortable pattern of early life failures (pattern F).

5 Is maintenance really necessary?

The lessons so far are:

- Failure does not stop happening when a part is in storage
- The failure modes experienced by equipment in a warehouse can be very different from those that occur during its normal working life
- The consequences of failure in store are usually very different from those during operation
- Most failures in store are hidden
- The part's accumulated hidden failures become painfully evident when it is first used

- Even a low rate of failures in store can translate into a very high chance that the part fails immediately after installation

These failures need to be managed: as we have seen already, even a very small rate of in-store failures can translate into an embarrassing chance of failure when the part is needed.

The reliability required from critical spares varies, but it is routinely above 95% and very often more than 99%. To meet this target on installation, a part is on the shelf for six months can have a failure rate of no more than 2% per year, or a mean time between failures of about 50 years. If the part's intrinsic, left-alone-by-itself failure rate is any higher, then maintenance in stores is not optional.

6 A Maintenance Programme

We saw in the last section that acceptable reliability can only be achieved by managing the inevitable risk of failure in the warehouse. Unless a part is guaranteed to be delivered in perfect condition and it has a spectacularly low failure rate, or the consequences of failure are trivial, these failure modes need to be managed through a maintenance programme. The obvious and most robust way to construct the maintenance schedules is by using Reliability-centred Maintenance (RCM).

This note isn't the right place for an RCM training course, but the basic principles applied to maintenance in stores are simple.

- A maintenance task should only be done if it is *technically feasible* and *worth doing*
- A task is *worth doing* if it achieves your organisation's goals. In the case of maintenance in stores, the task has to reduce the risk of failure on installation to a tolerable level. If it doesn't do that, the task is not worth considering.
- A task is *technically feasible* if it is capable of preventing, predicting or detecting a failure mode

¹ If you work with failure distributions you can see that I have taken some liberties when calculating failure probabilities in this paper. The failure rates here are low enough that the differences from an exact calculation are not too significant.

RCM NOTES

Most failures in store are hidden, but that does not mean that the only periodic tests (failure-finding) should be considered. The table below

shows the types of task that could be applied to typical in-store failures.

Type of task	Manages these failures	Examples	Characteristics
Condition-based	Corrosion, oxidation, slowly developing loss of containment	Visual checks for corrosion Check for signs of fluid leaks	There must be a clear, detectable potential failure condition before the part has completely failed
Scheduled restoration (overhaul)	Dirt, dust, distortion of stationary components, binding of moving parts	Clean equipment and components. Move rotating components, stroke valves	There must be a known minimum age at which the failure occurs
Scheduled discard	Breakdown of lubricants and other fluids, degradation of seals and plastic components by oxygen or light	Replace expired lubricants and other fluids Replace seals Replace desiccant	There must be a known minimum life at which the failure occurs
Failure-finding	Most failures that cannot be managed by condition-based or scheduled tasks	Test to detect failure of electronic components Visual and functional detection of accidental damage	It must be possible to detect whether the failure mode has occurred.

RCM is used to select in-store maintenance tasks in exactly the same way as it is for operational equipment, but there are a few points to bear in mind.

The consequences are different

As we have seen already, it is very likely that the consequences of failure, even of the same failure mode, will be very different from those the consequences when the equipment is operational. You will probably find yourself in the *hidden failure* section of the decision diagram far more often.

Failure development times can be different

The interval between potential failure and functional failure, the P-F interval, will often be different. Where failure is subject to a specific life, the time taken to fail may differ substantially.

Failures that involve deterioration of materials such as corrosion or degradation of seals may happen more or less quickly than during

operation, depending on the materials used and the operational and storage environments.

Differences in P-F intervals and component life translate directly into changes to maintenance task intervals, so expect the in-store maintenance intervals to be different from those for the same parts in deployed equipment.

Task feasibility changes

Remember that a task that is easy to carry out in the field may be tricky when equipment is on the shelf. Visually checking for seal leaks could take seconds when a pump is operational, but the same task might be completely impractical when the same part is on the shelf. The opposite is also true: some tasks are far easier in the storeroom, particularly those that need close access to the equipment.

7 Failure-finding in Store

The RCM decision logic prioritises condition-based tasks, scheduled restoration and discard over failure-finding for a very good reason: these tasks either prevent or predict a functional failure,

so there should only be a very small chance that the organisation will be exposed to the consequences of an unplanned breakdown. On the other hand, selecting a failure-finding task allows the failure to happen. A regular failure-finding task controls the chance that a part works when it is needed, but there is always a finite chance that the device has failed. For that reason, failure-finding should only be considered if there is no applicable condition-based or life-based task.

How do you know how often to carry out a failure-finding task? The interval is calculated by applying the same formulae as for operational equipment, based on one of three criteria:

- Device availability
- Multiple failure rate
- Cost of failure-finding versus cost of multiple failures

The formulae used, and the assumptions behind them, can be found at

<http://www.numeratis.com/book/realconsequences>

Device availability

This is the simplest calculation to carry out, but you need to choose an appropriate level of availability. The availability is the chance that a part is working (i.e. has not failed in the store) when it is demanded.

The calculation may be easy, but it isn't obvious how to choose an availability level. One way to start is to find out the chance that a part will be available from the store when it is needed. You will need to have access to simple spares optimisation tools to calculate an approximate value. This value (95%, 99% or whatever) is the chance that a demand for stock will be met from stock without waiting. The number probably represents the absolute minimum availability you need for your hidden failure calculation. If you choose a lower value, then part failures in store will contribute more to downtime costs than part stockouts.

Multiple failure rate

You need three values to carry out this calculation.

- The mean time between failures of the part in store
- Demand rate: the mean time between demands for the part
- Tolerable mean time between multiple failures

Determining the tolerable mean time between multiple failures needs some thought. This figure represents how often, on average, a multiple failure is allowed to occur; in other words, how often a demand is made for a part, but the part is not working or stops working immediately due to failures that developed in the store.

Cost balance

This calculation tries to balance the cost of regular failure-finding tasks against the risked costs of multiple failures. It needs four values.

- The mean time between failures of the part in store
- The mean time between demands for the part
- The cost of a single failure-finding task
- The cost of a single multiple failure, where a part has failed in the store when it is needed

The calculation determines the failure-finding interval that minimises cost over a long period of time.

8 Conclusion

Spare parts don't stop failing just because they are in a warehouse, and the reality is that most of those failures are hidden. When the part is needed to carry out a repair, failures that have occurred in store become embarrassingly evident. RCM—and specifically its treatment of hidden failures—provides a robust framework for the development of an in-store maintenance programme.

Terms of use and Copyright

Neither the author nor the publisher accepts any responsibility for the application of the information and techniques presented in this document, nor for any errors or omissions. The reader should satisfy himself or herself of the correctness and applicability of the techniques described in this document, and bears full responsibility for the consequences of any application.

Copyright © 2013 numeratis.com. Licensed for personal use only under a Creative Commons Attribution-Noncommercial-No Derivatives 3.0 Unported Licence. You may use this work for non-commercial purposes only. You may copy and distribute this work in its entirety provided that it is attributed to the author in the same way as in the original document and includes the original Terms of Use and Copyright statements. You may not create derivative works based on this work. You may not copy or use the images within this work except when copying or distributing the entire work.