Failure-finding without FEAR

RCM Notes series

Dr Mark Horton, Numeratis.com, March 2011

Introduction

Your RCM analysis has been going well. Yes, the functions had caused some problems and you'd had long discussions about some safety issues, but since you started on the decision worksheet it has been plain sailing. Now you are through the primary functions and three secondary functions. The end is finally in sight and you can look forward to a well-earned break doing your paid job.

And then it happens.

> *OK, we're on function 5, functional failure A, failure mode 1. This is 'High temperature trip system fails'. Joe, will you read through the failure effects?*
>
> *Thanks, Joe. Let's start at the top. 'Will the loss of function caused by this failure mode on its own become evident to the operators under normal circumstances?'*
>
> *"No"*

It's about now you get that sinking feeling in your stomach. Can this really be happening when everything was going so well?

> *Let's just take another look at that. Wouldn't the operator see this failure happen?*
>
> *"No"*
>
> *Er… what about the maintainer? Fred, what do you think?*
>
> *"I don't see how we'd ever know until we got a high temperature in there."*
>
> *Fine. I suppose we're going down the blue column.*

It could be worse. There's a chance you will tackle it with condition monitoring. All isn't lost… yet.

> *We're going to look at the first question. Is it possible to define a detectable potential failure condition?*
>
> *"No"*
>
> *Let me put it another way. Before the trip fails, are there any warning signs that could tell you that the failure is going to happen?*
>
> *"No"*
>
> *None at all?*
>
> *"No, none at all."*
>
> *I suppose that's a "no" to the first question. Let's look at scheduled restoration.*

You knew that today was going to be like this. But there could be some hope left.

> *Is there an age at which there is a rapid increase in the conditional probability of failure?*
>
> *"No"*
>
> *Sure?*
>
> *"Yes."*
>
> *… I don't suppose scheduled discard will help?*
>
> *"No."*

Morale is hitting rock bottom. Now where did you put those facilitator course notes? Will they notice if you slip out and feel sick for a few minutes?

> *So now we're down at failure-finding. Joe, did you say you were going to bring in the doughnuts today? I think that now would be a good time to take a break…*

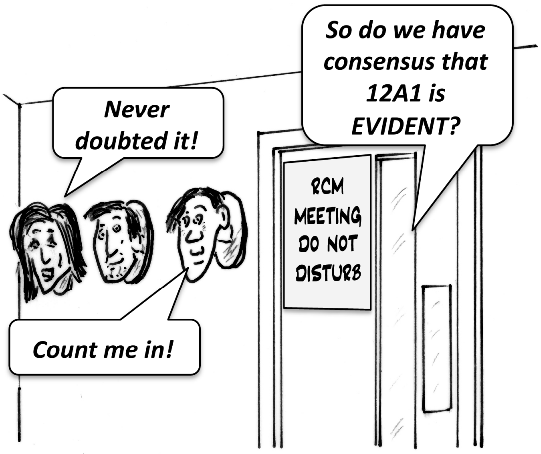
So do we have consensus that 12A1 is EVIDENT?
Never doubted it!
Count me in!
RCM MEETING DO NOT DISTURB

No one ever said that RCM facilitation would be stress-free, but over the years we have found facilitators' stress levels rising (and morale falling) more in the Hidden column than anywhere else in the decision diagram.

The objective of this workshop is to reduce your stress and give you an easier life as a facilitator, at the same time as looking at the issues raised by the failure-finding question. We're going to assume that you know the failure-finding formulae—or you can look them up—and we will focus on

- How you get the information that you need to do the calculations
- How to make the facilitation of hidden failures as trouble-free as possible

We'll start at the top of the hidden column and look at some typical problem areas and solutions, then look at the failure-finding question itself.

Surviving Failure-Finding

These notes are only a summary of the approach to failure-finding calculations. If you are in any doubt at all, please refer to your RCM trainer or mentor.

1 Consider on-condition and fixed time maintenance first

Don't jump immediately to failure-finding for a hidden failure. It is better, if possible, to deal with the failure through condition monitoring, scheduled restoration or scheduled discard because these prevent the failure rather than allowing the device to fail. Only start to ask the failure-finding question if the group has already answered *No* to the first three questions.

2 Is it possible to check whether the device has failed?

Ask this question before starting the failure-finding calculation. If the answer is *No*, the answer to the failure-finding question is *No* and you do not need to carry out the calculation at all.

If the answer is *Yes*, take a deep breath at this point.

3 How complex is the system?

The simple failure-finding formulae are (unsurprisingly) applicable only to simple systems. These are typically a single protective system protecting against a single demand. The protective system may consist of identical parallel devices, but don't try to use the simple formulae for anything more complex: consult your local RCM or risk analysis expert.

4 Risk or Economic?

Whether you need to ask this question depends on the variety of RCM decision diagram you are using. Some varieties split hidden failures into Hidden/Safety, Hidden/ Operational and Hidden/Non-Operational. Others, including RCM 2, have only a single Hidden column, and you need to decide whether the failure-finding calculation is based on economics or risk.

Describe the multiple failure to the group using the names of the demand and protective device:

"The multiple failure is that the lubrication system low low level trip has failed and the oil level drops. The pump does not shut down and it could be damaged so badly that it has to be replaced."

Ask whether there are any safety or environmental consequences associated with the multiple failure. Select the formula appropriate to the system you want to analyse. The simple (single device) formulae are shown below.

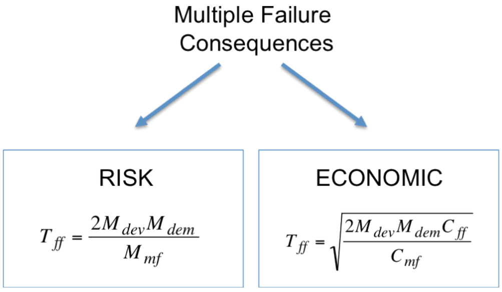

$$T_{ff} = \frac{2M_{dev}M_{dem}}{M_{mf}}$$

$$T_{ff} = \sqrt{\frac{2M_{dev}M_{dem}C_{ff}}{C_{mf}}}$$

5 Write down the terms

Write down all the terms in the formula you have chosen. You don't have to write these on a flipchart yet, but you need to know exactly what information you need so that you don't miss essential data or waste the group's time by trying to find information that isn't needed at all.

Simple Risk Formula	Simple Economic Formula
M_{dev}	M_{dev}
M_{dem}	M_{dem}
M_{mf}	C_{ff}
	C_{mf}
T_{ff}	T_{ff}

6 Translate the terms

Write the terms at the left of a flipchart page.

First describe each term in terms of the failure-finding formulae.

Then translate each item into the terms of the exact system you are analysing and write the description next to the mathematical symbol as shown below.

Simple Risk Formula

Term	What it means	Value
M_{dev}	Mean time between failures of the relief valve (how often, on average, it jams shut)	
M_{dem}	How often on average we call on the relief valve because the boiler goes overpressure	
M_{mf}	How often we are willing to accept that boiler blows up because the relief valve is jammed closed	
T_{ff}	How often we will test the relief valve	

Simple Economic Formula

Term	What it means	Value
M_{dev}	Mean time between failures of the motor overload trip (how often, on average, it would fail to a state in which it cannot trip the motor)	
M_{dem}	How often on average we call on the motor overload because the motor is stalled	
C_{ff}	How much it costs to check that the overload is operational every time we carry out the test	
C_{mf}	How much it would cost if the multiple failure occurred; i.e. that the motor is stalled, the overload does not trip it and the motor burns out	
T_{ff}	How often we will test the motor overload trip	

7 Get the Values

Now you know what the terms in the equations refer to, you need to fill in the values. This is probably the most difficult part of the whole process. Focus on getting the information and, more importantly, recognising what you don't know. The steps shown here are for guidance only, and you should consult your own company's safety analysis procedures where appropriate.

7.1 M_{dev}

You are trying to find out the failure rate of each individual protective device (alarm, trip, overload, relief valve) if it were left by itself without maintenance.

You can start by asking if the device is checked at the moment, and if it has ever been found in a failed state.

1 *"Do you check this alarm/trip/relief valve at the moment?"*

2 *"If you do, have you ever found it failed when you did a check?"*

3 *"How many times have you found it failed? Over what period?"*

Then calculate M_{dev}

$$M_{dev} = \frac{Period}{Number\ of\ \ failures}$$

If this does not work, ask:

1 *"Are there any other alarms/trips/relief valves like this one on site?"*

2 *"If there are, have you ever found any in a failed state?"*

3 *"If so, how many times over what period?"*

Then calculate M_{dev}

$$M_{dev} = \frac{Period \times Number\ of\ \ devices}{Number\ of\ \ failures}$$

If that doesn't produce any information you can use, ask

1 *"Is there anywhere (manufacturer, generic data...) where we could get this information?"*

If the answer is *"No"*, the workshop will give some guidance on estimating the possible values.

7.2 M_{dem}

You are trying to find out how often the protective device has to operate for real (not on test).

You can start by asking if the protective device has ever been needed.

1 *"Have you ever activated this alarm/trip/relief valve?"*

2 *"How many times have you done it? Over what period?"*

Then calculate M_{dem}

$$M_{dem} = \frac{Period}{Number\ of\ activations}$$

If this does not work, ask:

1 *"Have you ever had any near misses which might have needed the alarm/trip/relief valve?"*

2 *"If there have been, what is the chance of a near miss turning into an incident?"*

Then calculate M_{dem}

$$M_{dem} = \frac{Period}{Number\ of\ near\ misses \times Chance\ of\ incident}$$

If this does not work, ask:

1 *"Are there any other alarms/trips/relief valves on systems like this one on site?"*

2 *"If there are, have you ever activated those?"*

3 *"If so, how many times over what period?"*

Then calculate M_{dem}

$$M_{dem} = \frac{Period \times Number\ of\ systems}{Number\ of\ activations}$$

If that doesn't produce any information you can use, ask

1 *"Is there anywhere (other users of similar systems...) where we could get this information?"*

7.3 M_{mf}

There are no ways to estimate this figure.

Ask

"How often would we be willing to have the boiler go overpressure with the relief valve jammed closed so that the boiler explodes?"

1 Ensure that you consult management if the multiple failure has safety or environmental consequences

2 Remember that there may be many hidden failure modes on site which ultimately have safety or environmental consequences. If there are 100 failure modes like this one on site and M_{mf} for each is 10000 years, the mean time between multiple failures for the whole site is 100 years.

If no one is willing or able to set a figure, the most important rule is to *err on the side of safety*. The following approach may be useful, but it should be used with caution to give you an idea of whether failure-finding is going to be appropriate. Be aware that most individuals aren't used to dealing with levels of risk in their jobs or personal lives, and our assessments as human beings can be wildly inaccurate when the probability of an event occurring is low.

1 Find M_{dev} and M_{dem} as usual

2 Assume a practical failure-finding interval, perhaps based on the current interval

3 *Calculate* M_{mf} (this applies only to a single protective device)

$$M_{mf} = \frac{2M_{dev}M_{dem}}{FFI}$$

4 Ask

"Is this mean time between multiple failures acceptable?"

If it is very short and the failure-finding interval cannot be substantially reduced, redesign may be appropriate.

8 Plug the Numbers into the Formula

Complete the table with the figures, plug all the figures into the formula and calculate the failure-finding interval.

Remember to express all the times in the same units (usually years).

Term	What it means	Value
M_{dev}	Mean time between failures of the relief valve (how often, on average, it jams shut)	70 Years
M_{dem}	How often on average we call on the relief valve because the boiler goes overpressure	100 Years
M_{mf}	How often we are willing to accept that boiler blows up because the relief valve is jammed closed	100000 Years
T_{ff}	How often we will have to test the relief valve	0.14 Years (6 weeks)

Term	What it means	Value
M_{dev}	Mean time between failures of the motor overload trip (how often, on average, it would fail to a state in which it cannot trip the motor)	100 Years
M_{dem}	How often on average we call on the motor overload because the motor is stalled	25 Years
C_{ff}	How much it costs to check that the overload is operational every time we carry out the test	£20
C_{mf}	How much it would cost if the multiple failure occurred; i.e. that the motor is stalled, the overload does not trip it and the motor burns out	£3500
T_{ff}	How often we will test the motor overload trip	5 Years

9 *Checks and Balances*

It isn't all over just yet. Most of the formulae used for failure-finding calculations are valid only for a range of values. If you're outside their range of validity, your failure-finding values may be incorrect. Worse, failure-finding might not be the right option at all for this failure mode.

9.1 Formula Validity

Work out the unavailability for a single protective device. Use the same formula even if you are working with parallel redundant protective devices, because the approximation applies at the level of an individual protective device.

$$U = \frac{T_{ff}}{2M_{dev}}$$

Now check if U is greater than about 0.05 (5%). If it is greater than 0.05, the formula is outside its range of validity.

9.2 Task Feasibility (disabling the device)

Check the availability figure that you calculated above.

$$U = \frac{T_{ff}}{2M_{dev}}$$

If it is very low—exactly how low depends on the task you have in mind—you must seriously question whether the task can be done well enough to guarantee that level of unavailability.

Example

A pressure switch is used to shut down a process if the pressure in a reaction vessel rises above 10 000 kPa. If the pressure switch failed to operate and the vessel pressure rose violently, the vessel could explode and cause a reportable environmental incident. Although no one is usually present in the area, it is possible that maintenance or operations personnel could be injured or even killed in the incident.

After some discussion, the RCM group decides to calculate a failure-finding interval for the switch based on the following data:

Term	What it means	Value
M_{dev}	Mean time between failures of the pressure switch (how often, on average, it fails closed)	250 Years
M_{dem}	How often on average we call on the pressure switch because the reactor vessel goes overpressure	10 Years
M_{mf}	How often we are willing to accept that reaction vessel blows up because the pressure switch has failed	50000 Years
T_{ff}	How often we will have to test the pressure switch to achieve (based on risk formula)	5 Weeks

The task the group has in mind is to isolate the pressure switch at a local block valve and attach a small pump to pump up the isolated leg to the trip pressure.

The facilitator calculates the unavailability $U = T_{ff}/(2M_{dev}) = 0.0002$ (0.02%).

The group thinks that the chance of leaving the pressure switch disabled after the test is not less than 0.02% (one in 5000) and might be as high as one in 1000, so that task as envisaged is not feasible.

If you find that the required availability is very high, the options available to you are:

- Look for some other, more foolproof way to test the device. This might involve testing a whole instrumentation loop rather than isolating sections of it to carry out individual tests.
- Make the test instructions very detailed and incorporate double-checking at the end of the procedure (*"...After completing the test, the supervisor is to check independently that the block valve to the pressure switch has been opened. He/she must sign the job sheet to confirm that the check has been made."*)
- Answer *No* to failure-finding and consider redesign options

9.3 $M_{dem} >> T_{ff}$

Ensure M_{dem} is much bigger than T_{ff}. If it is less than T_{ff} or about the same value, failure-finding is likely to be ineffective since the device is being operated just about as often as it is being tested.

Example

An overhoist switch protects a crane from being raised too far and possibly dropping its load. The operator currently hits the limit about once per day. Therefore failure-finding at an interval greater than one day is unlikely to have much effect on the availability of the device.

This condition usually indicates **alarm** or **trip abuse**: the system is poorly controlled and the alarms/trips are being used as control systems, not as emergency systems.

The options here are *redesign* or *no scheduled maintenance*. If *no scheduled maintenance* is chosen, the mean time between multiple failures for a single protective device is very roughly

$$M_{mf} \approx M_{dev} + M_{dem}$$

for M_{dem} very much less than M_{dev}.

Is this mean time between multiple failures acceptable?

Learning to Love Failure-Finding

No one ever said that failure-finding would be easy, but by knowing the basics and applying them consistently, you can not only reduce your stress level but also keep the motivation and attention of your review group. In time, you may even get to enjoy dealing with hidden failures!

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.