4 Failure-Finding Basics

4.1 Introduction

This chapter builds the foundations that you will need to apply failure-finding and other failure management policies to real equipment.

The techniques presented in this book are like tools in a toolbox. Before using them, you need to be able to understand the terminology and to identify the protective device, the demand, and the ultimate multiple failure. Even if you already have some background in risk analysis, risk-based inspection or Reliability-centred Maintenance, you should spend some time becoming familiar with the terminology used in the following chapters.

4.2 Protective Devices and Systems

The terms *protective device* and *protective system* are used interchangeably in this book.

A protective device is intended to operate if an initiating event or trigger event occurs. In general the term "protective device" is used for a small, self-contained component such as a sensor or a relief valve, while "protective system" is applied to a whole item of equipment such as a fire alarm. The terms are often used interchangeably in this book, and there is not usually any significance in the use of "device" rather than "system".

Examples of protective systems are listed below.

Protective System
Fire alarm
Pressure relief valve
Pump motor trip
Car anti-lock braking system (ABS)
Hospital emergency generator

4.3 Demand and Initiating Event

The protective device operates when a *demand* is placed on it by an *initiating event* or *trigger event*. These three terms are used interchangeably.

Examples of typical demands on protective systems are listed below.

Protective System	Demand (initiating event)
Fire alarm	A fire breaks out
Pressure relief valve	Steam boiler overpressure
Pump motor trip	The pump motor stalls
Car anti-lock braking system (ABS)	Need to brake in an emergency or in slippery conditions
Hospital emergency generator	Main electric power supply failure

Protected Function

The term *protected function* is used by Moubray (1997) and in other published work derived from RCM 2. This book avoids using the term for a number of reasons.

- "Demand" and "event" are far more widely accepted, and they clearly describe the relationship between the protective device and the events that should cause it to operate
- The terms "protected function" and "protective device" are so similar that they often cause confusion
- It is the *failure of the protected function* that actually places a demand on the protective device
- It is sometimes unclear what the protected function actually is

If the protective system is a backup system such as a standby water pump, it is obvious that the protected function is something like this: "To pump water at a specified rate", a function that is probably part of the RCM analysis. It is far less clear if the device is a fire alarm, where the function could be "Not to catch fire", which would probably not appear in the analysis. Overall, the term *protected function* has been avoided to improve clarity.

4.4 Multiple Failure

The multiple failure is what happens if the demand occurs while the protective system is in a failed state. The effects of the multiple failure need to be recorded clearly so that your can set up a consistent maintenance schedule for the protective system.

Failure Effects and Consequences

Before beginning to analyse a protective system, ensure that the following components are clearly identified.

- The protective system
- The demand or initiating event
- The multiple failure

Do not be tempted to continue with the analysis until you can clearly define each of the above elements. If you are facilitating an RCM review group, consider writing them down so that no one is in any doubt.

The following table shows some examples of protective systems, the associated demands and a definition of the multiple failure in each case.

Protective System	Demand	Multiple Failure
Fire alarm	Fire	An undetected fire occurs, resulting in increased risk of death, injury and physical damage.
Pressure relief valve	Steam boiler overpressure	Excess steam pressure is not relieved and the boiler explodes resulting in death and injury of personnel.
Pump motor trip	Motor stall	The motor stalls and burns out.
Car anti-lock braking system (ABS)	Need to brake in an emergency	ABS does not operate when brakes are applied in an emergency, and the vehicle skids out of control.
Hospital emergency generator	Main power supply failure	Emergency generator does not start during a power outage.

4.5 Failure Modes

In Chapter 2 we saw that one protective device can fail in a number of ways; in other words, it displays a number of *failure modes*. The primary function of the device may be hidden, but that does not mean that all of its possible failure modes are also hidden.

Protective devices can fail in two distinct ways: fail to operate when required, and to operate when there is no demand (spurious operation). A device can be subject to both hidden and evident failure modes. *Examples*.

4.6 Availability

Beware: "Availability" is a deceptively simple word. A protective device is *available* if it is capable of performing its function if a demand occurs. If it is incapable of correct operation, it is *unavailable*. From this point of view, a protective system is either available or it is not, so its availability is either 100% or 0%. In the real world, availability could hardly be a simpler concept.

Mathematicians, statisticians and reliability engineers learn over a period of many years' training to make simple ideas far more complex. To a reliability engineer, the availability of a protective device could be 0%, or 100%, or any number in between. To see how this picture differs from the simple all-or-nothing, 100% or 0% picture of availability, consider the following question.

"Did the fire alarm operate when we had that electrical fire last week?"

This is a simple question, and the answer is equally simple: either it worked or it didn't. The question could be rephrased in availability terms like this:

"What was the availability of the fire alarm when we had that electrical fire last week?"

The answer is either 100% or 0%, not 80% or 99.5% or 5%. It worked or it didn't.

Now look at a different question.

"If a fire were to occur now, would the fire alarm be capable of detecting it and annunciating an alarm?"

The truthful answer to this question is that we have no idea. In availability terms, the question is:

"What is the availability of the fire alarm now?"

There are two different ways in which we could try answer this question. Since the real world answer to the question is either 100% or 0%, we could start a fire (or preferably simulate one) and see whether the fire alarm operates. If it does, it was available; if it doesn't, it was unavailable. Although that gives us a definite answer, it's of no real use to us. Truthfully we don't want to know whether the alarm works now; we are far more concerned about whether it would operate when no one is around to test it, perhaps in the middle of the night. What we want to know is:

"What is the chance that the fire alarm would work if a fire occurred?"

An analysis of the system and its maintenance (perhaps by a reliability engineer) might be able to tell us the *probability* that the alarm would work correctly if a fire occurred. Although the "all-or-nothing" picture of availability represents what happens when a fire occurs, this probability is of far more use to us. It tells us how likely our protective systems are to operate when they are needed. The probability of operation is a number between 0% and 100% and it is known as the availability.

The probability of the alarm working correctly depends on a number of factors that we will investigate in the following sections.

In order to be effective, a protective system not only needs to exist, it needs to be available when it is required. For example, a simple fire alarm system could be unavailable for a number of reasons when a fire occurs.

- A component has failed in such a way that it is unable to detect a fire and annunciate an alarm
- The system has recently tested in a way that involves disabling part of the system during the test, and the technician forgot to enable the system afterwards
- The system's power supply has failed and no backup power supply is available

Any one of these failures is sufficient to ensure that the fire alarm's function is unavailable when a fire occurs. While it is possible to influence a system's availability through scheduled testing, failure of its components is only one root cause of unavailability. The simple example above demonstrates that unavailability may also arise from human intervention (testing) and external factors (the power supply) and even its design. When analysing a protective system, ensure that you understand and take into account all the factors that might disable it, not just those which maintenance can influence.

For discussion

The fire alarm in this example is a simple system. Most commercial systems incorporate a battery back-up power supply so that they can operate for extended periods without mains power; the alarm may also signal its monitoring centre when power supply problems occur.

What additional maintenance requirements could arise because of the increased complexity of a fire alarm which includes a back-up power supply and signalling, compared with the maintenance of a simple alarm?

4.7 Availability: a Practical Example

What does availability mean for a real system? How does availability depend on the maintenance policy chosen for the protective system?

To answer these questions we will calculate the availability of a fire alarm system during one calendar year. The alarm is known to be working at the start of the year, but it fails a few moments after midnight in the morning of 1 April. In this first example, the alarm is not checked again until the end of the year. What is its availability over the year?

Let us be clear about the sense of the word "availability" in this section. We do not mean, for example, "Does the alarm function when a cigarette starts a fire on 18 July?" The availability that we want to determine is the probability that the alarm would operate if a fire occurred on a randomly chosen day during the year. We assume that no fires actually occur during the year.

In this first example, the alarm system is operational from 1 January to 31 March. It fails, but the failure is hidden because no fire occurs. The failure is discovered at the end of the year and the alarm system repaired.

What is the availability of the alarm over the year? Because we have the benefit of perfect knowledge, we know that the alarm was operational from 1 January to 31 March, or 90 days. The system availability is therefore

$$\frac{90}{365} = 24.7\%$$

In the year we have chosen, the availability of the alarm system is poor. What effect can a different maintenance policy have on the availability achieved?

What is the availability if it is tested on 1 January and 1 July rather than just once per year? The system is now operational from 1 January to 31 March when it fails. Since the failure is hidden, it remains in a failed state until it is tested on 1 July. It is tested, found to be failed and repaired. For the sake of simplicity we assume that no further failures occur during the remainder of the year. What is the overall availability achieved?

The system is again available for 90 days to 31 March; from 1 April to 30 June it is unavailable (91 days); it is repaired on 1 July and is operational for the rest of the year (184 days). It is therefore available for 274 out of 365 days, or 75% of the time.

The table above summaries the availability achived for a range of task intervals.

Test interval	Operational Days	Non-operational days	Availability
1 Year	90	275	25%
6 Months	274	91	75%
1 Month	335	30	92%
1 Week	358	7	98.1%
1 Day	364	1	99.7%

The graph below shows how the alarm availability changes as the testing interval is increased from 1 day to 1 year. It may be surprising that the graph is not a smooth curve, but remember that we have made a number of assumptions. First, the device is checked on 1 January. We assume that there is a single failure on 1 April, when in a real situation we would have no idea when the failure might happen, because the failure is random and hidden. Finally, availability is calculated over the year to the end of 31 December, not over a long—or possibly infinite—period, as it might be in the models that we will use shortly.

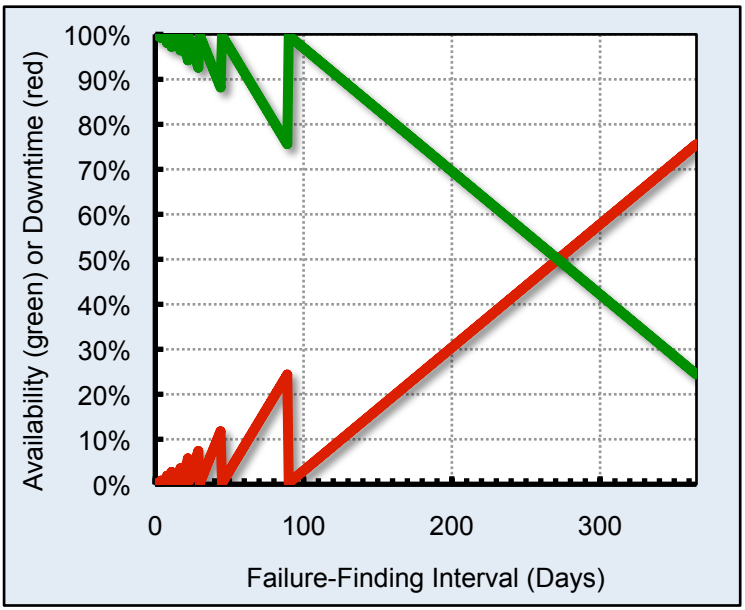

Availability (green) and downtime (red) for failure-finding intervals from one day to one year, assuming that the device device has been checked on 1 January and that it fails on 1 April

The availability achieved has peaks and troughs depending on how close to 1 April the task is carried out. So if the task interval is 89 days, the first task after 1 January just misses the failure, so the failure is not found until the second task, resulting in downtime of nearly 25%. If the task interval is 91 days, the first testing task catches the failure, and downtime is only a single day over the year.

In a similar way, the table below examines the effect of increasing the testing frequency.

This exercise demonstrates that there is a relationship between availability and testing frequency: ignoring for a second the peaks and troughs shown on the graph, the protective device spends less time in an undetected failed state if it is tested more often, and so a higher overall availability is achieved.

As has already been pointed out, this section is in some ways a fraud because real life is very different from the simple example above.

First, the assumption that the alarm fails on 1 April is unrealistic. If we knew that the device would fail on 1 April, we would intervene in some way to provide continuous alternative protection or to repair the alarm as soon as possible. There are two reasons why this assumption is unrealistic. First, unless the device has a very well-defined lifetime, we have no idea exactly when it will fail. Second, because the failure is hidden, there is no way for us to know that the failure has occurred except to test it.

Second, we have assumed that we can "re-run" the same year's history with different task intervals, certain that the failure will occur on 1 April every time. If failures of the protective device occur at random, then history is absolutely no guide to the future, and no one year will be like the one before or the next.

Unrealistic as it is, the example does demonstrate one fundamental principle very clearly: that protective device availability is not a property that is fixed by the manufacturer and over which we have no influence. In summary,

If the protective device works when it is first installed, its availability is entirely controlled by our maintenance policy.

This is why it is vital to pay close attention to the failure-finding interval and to the way in which the task is carried out. Calculating the failure-finding interval is the core subject matter of section 2.

4.8 Key Points and Review

A protective device is designed to initiate a response if an unusual condition (the demand) occurs.

The protective device is usually designed so that, if it performs correctly, it reduces or eliminates the consequences of the demand.

A multiple failure occurs if a demand arises when the protective device is in a failed state or disabled in some other way (a fire occurs but the fire alarm is broken or turned off, so people are at increased risk of death or serious injury).

Protective devices can fail to operate when required (the multiple failure); they can also operate when they are not required (spurious operation).

Choosing failure-finding as a maintenance policy for a protective device means that the device can be in a failed state for an extended period, so a multiple failure could occur.

In the simple model presented in this chapter, the availability of a device can be increased by checking its operation more frequently.

The rate at which multiple failures occur can be managed in two ways: by increasing the availability of the protective device (for example, by checking it more frequently); and by reducing the demand rate (possibly by maintenance on or redesign of the system which causes the demands).

The objective of any management policy is to reduce the chance of a multiple failure to a tolerable level.