3 Managing Hidden Failures

3.1 Introduction

Hidden failures need to be managed because of the severity of the consequences of a multiple failure. Managing hidden failures poses two specific challenges. First, and by definition, there are no observable effects when a hidden failure occurs by itself. This is precisely what makes the failure hidden rather than evident. Secondly, many of the devices used in protective systems rely on electronics and other technologies that predominantly fail at random, with no predictable pattern or age of failure. These two factors together appear to make the management of hidden failures an impossibility.

A management policy that focuses on the *effects* of a hidden failure is doomed, simply because there are no effects to manage. The key to management of hidden failures is to focus on the characteristics of the failure itself rather than on its effects. What are the characteristics that could provide the basis of a successful failure management policy? This chapter examines the factors that affect the selection of maintenance tasks in general, and with an emphasis on hidden failures in particular.

3.2 One Part, Several Failure Modes

One part can have several failure modes. Some of them may be hidden, others evident. A failure maintenance policy is needed for each failure mode, not just one policy for the whole part.

3.3 Scheduled Overhaul and Discard

Many failure modes have a characteristic life. Their life is not a point at which all failures will occur, but it is a time when the probability of failure starts to increase rapidly. If the part remains in service, it becomes more and more likely to fail.

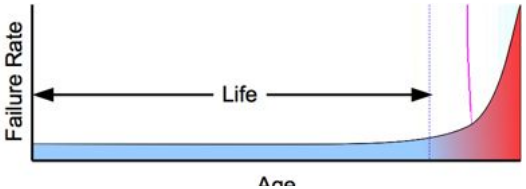

An age-related failure pattern

In the pattern illustrated above, there is a small, roughly constant chance that the failure occurs at any time after the part is installed. Later on, the chance of failure begins to increase rapidly; this point is marked as the item's life.

The obvious response to items that have an identifiable life is to replace them before that life is reached. This type of task is known as *scheduled discard*, *scheduled replacement* or *lifed* task.

An item's life is almost never known exactly unless formal reliability trials have been carried out. In general an estimate is made of the likely minimum life, and the replacement task is scheduled before that life is reached. Life is not always measured in terms of calendar time: it may be expressed in run hours or some other measurement of the part's usage. The common characteristic of all scheduled discard tasks is that they are carried out at fixed intervals.

This is obvious when considering evident failures, since these are the drivers for cyclic replacement of components.

Examples of "lifed" items include the following.

Component	Failure Mode(s) driving replacement cycle
Vehicle tyre	Tread wear Material degradation
Pipe	Erosion by impact of particles in fluid Corrosion by fluid External corrosion
Pump impeller	Erosion
Aircraft wings	Fatigue

Some components are subject to several "lifed" failure modes. This may lead to an "either/or" maintenance policy. For example, in the case of tyres, the material of which they are made wears off (lifed failure mode 1) and also degrades over time (lifed failure mode 2). This leads to a task that could be written as follows.

"Replace tyres every 50000 km travelled or every 5 years"

If a pipe would fail after five years as the result of internal erosion, but after 15 years because of external corrosion, it might be replaced every five years. However, as we shall see below, it is often possible to devise a maintenance policy for some items which is both safer and extends a component's useful life compared with fixed interval replacement.

An important point to take from this section is that the "life" which drives any task is that of the *failure mode*, not that of the item or part. As with the tyre, a single part may be subject to several failure modes, each of which has a different characteristic life. A separate task is needed to manage each failure mode, and the tasks are then combined when the maintenance schedule is constructed.

3.4 Condition-based Maintenance

Research in civil aviation during the 1960s and 1970s revealed a fundamental problem with maintenance management that relies on scheduled replacement: most failure modes do not have a predictable life. Of the items studied by United Airlines, fewer than 11% had failure patterns for which lifed tasks would have been a plausible management strategy (Nowlan and Heap, 1978).

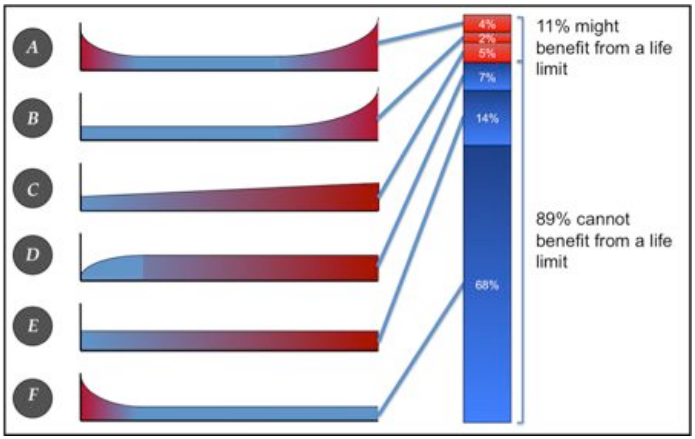

Equipment failure patterns from Nowlan and Heap (1978)

How, then, could the remaining 89% of failures be managed?

Although relatively few failure modes have a definite characteristic life, many failures give some warning that they are about to happen. The length of the warning period may vary widely, from seconds to months or years, but often it is long enough to prevent the failure from occurring, or at least to reduce or eliminate the consequences of failure.

The diagram below illustrates the failure development process from a point where the equipment is running acceptably through to the point of failure.

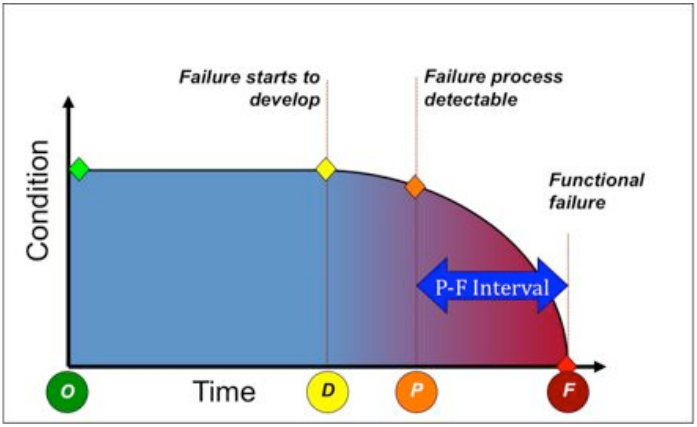

Failure development process through to functional failure

Initially the equipment is operating acceptably (point O above). At some point *D*, not necessarily related in any way to the equipment's age, the failure begins to develop; it may not be possible to detect any symptoms immediately, but if nothing is done, the item will deteriorate all the way to functional failure (point *F*). Between points *D* and *F*, it becomes possible to detect the deterioration, perhaps by sight, by sound, or by using some form of sensor. This is the point of potential failure (P), sometimes called incipient failure.

The example below demonstrates how this might apply to failure of an engine's piston rings.

Stage	Description
O	Engine running normally
D	Piston ring begins to wear. Metal particles are present in the oil but are not detectable by normal oil sampling techniques.
P	Detectable debris present in the oil. Oil leaks as wear increases. Loss of compression.
F	Engine is unable to sustain the required load. Serious oil leakage. Severe engine damage is possible.

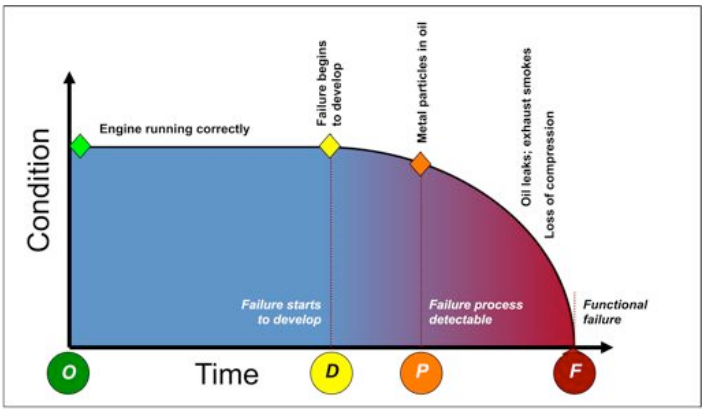

Failure development and potential failure conditions

A failing piston ring shows more than one warning sign before the engine stops: first, the presence of small metal particles in the engine oil, then oil leaks and lack of compression. Any of these symptoms might be used as a potential failure condition. The key for managing the failure is the interval between a detectable symptom of failure (particles, oil loss and so on) and final functional failure (the engine stops). This interval is known as the P-F interval.

Which symptom is chosen as the potential failure condition depends on the expected interval between the P-F interval and how much notice is needed in order to schedule appropriate maintenance or to mitigate the consequences of failure. Depending on how the engine is used, oil sampling might have a P-F interval of weeks or even months; compression testing or looking for a smoky exhaust might give days' or weeks' notice of failure.

When a potential failure condition has been identified, the condition monitoring task that manages the failure can be written. Because P-F intervals are often not known with any certainty, it is common to schedule the condition monitoring task to take place at half the P-F interval, although there is no absolutely definite rule. If we assume that the minimum P-F interval for oil sampling is eight weeks, then the task chosen to manage the failure might be written like this.

> "Take an oil sample every four weeks [i.e. half the P-F interval] and send sample for analysis. If analysis indicates piston ring wear, schedule replacement of all rings or substitution of engine."

3.5 Failure-Finding

Failures that occur after a definite life can be managed by scheduled component replacement; failures that provide some form of warning or potential failure can be managed by condition monitoring. What can be done if the failure has no life and no identifiable potential failure?

If we were trying to manage an evident failure mode, the answer would be simple: maintenance could do nothing. If the consequences of doing no maintenance were unacceptable, our only remaining option would be to redesign the equipment or to change the way in which we use it.

The situation is different if we are dealing with a hidden failure, simply because in normal circumstances we do not know that a failure has occurred. If we do nothing at all, the protective system will fail, then remain in a failed state until whatever event it is supposed to protect us against happens. And then, of course, it will do nothing because it has already failed.

The difference between hidden and evident failures gives us an alternative to doing nothing. Although we may not be able to predict that a failure is going to occur, we can detect it once it has happened. This is not going to prevent the failure, but we may be able to find out that the protective system has failed before its failure has any consequences.

For example, a fire alarm consists of many components, the failure of any one of which could prevent it from working if a fire started. Very few of these failures have a definite life or give any warning of failure before they happen. However, we can test the alarm, perhaps by pressing a "test" button, or by simulating a fire, and repair the alarm if it does not operate. If the test shows that the alarm is not working, then it can be repaired so that it is capable of detecting a real fire.

This testing or checking task, usually carried out at regular intervals, is called a *failure-finding* task. There is a very important distinction between failure-finding tasks and the scheduled discard, scheduled refurbishment and condition monitoring tasks discussed above. If a scheduled discard, scheduled refurbishment or condition monitoring task works as intended, the failure that it is managing never happens. Failure-finding tasks are different because they potentially allow the protective system to fail. More than that, if the protective system has failed, it remains in a failed state until the next failure-finding task is carried out.

Why does this matter? Because if the protective system fails, there is no protection *at all* until the next failure-finding task is carried out and reveals the failure. So if the fire alarm fails, it is incapable of protecting us from the consequences of a fire until it is next tested.

So failure-finding tasks are fundamentally different from the range of scheduled tasks that can also be applied to evident failures because of the possible gap between failure of the protective system and the next failure-finding task. The gap means that however often we check the protective system, there is always a small probability that it could be in a failed state when the event against which it is protecting us occurs.

3.6 Do nothing

To do nothing might appear to be an unlikely failure management strategy, but it is a logical choice under certain conditions.

- The failure has no safety or environmental consequences, and
- There is no applicable preventive task, or
- The cost of carrying out any applicable preventive task is more than the cost of allowing the failure to happen

Doing nothing, or more formally *no scheduled maintenance*, may therefore be a positive decision based on the reliability and cost data that are available.

Doing nothing is an unusual decision for hidden failures, because if the design of equipment requires a protective device, it is likely that maintenance is required to ensure that it is available when needed. Section 3.8 below describes the circumstances under which failure-finding may not be applicable to hidden failures.

3.7 Redesign Options

While redesign is not strictly part of maintenance management, it is an important aspect of failure management. If no maintenance task can prevent, predict or detect the failure, and doing nothing is not an acceptable option, the final choice is to redesign the equipment. "Redesign" does not necessarily mean a high cost, physical redesign; in practice, redesign may mean changing the way in which equipment is used or changing operating instructions.

Redesign as a management strategy for hidden failures is discussed fully in a later chapter. Examples might include the following.

Redesign Strategy	Examples
Make the hidden failure evident	An uninterruptible power supply checks itself three times per day and raises an alarm if its battery has failed Add a circuit to check the continuity of normally-off incandescent warning lights
Add more protection	Add a second relief valve to a vessel that has a single valve Add a tank ultimate high-level shutdown switch in addition to its existing alarm system
Improve the reliability of the protective system	Upgrade a tank's ultimate level switch to a new model that prevents the entry of dust, dirt and product
Reduce the rate of initiating events	Improve a crane's control system and retrain operators to reduce the rate of demand on the overhoist protection switch

3.8 When is Failure-Finding not Feasible?

If you believe that every hidden failure can be successfully managed by carrying out some form of scheduled test, then think again. Failure-finding is only a practical management policy if the following conditions are met.

- It is possible to check whether the protective system has failed, and
- It is practical to carry out the failure-finding task at the required interval

Both points probably seem obvious, but you need to be aware of some important issues that have to be considered.

When is it impossible to check whether a protective device has failed? Devices that cannot be tested without destroying them belong in this category. They include fuses, rupture discs, shear pins and automobile air bags among others. None of these devices can be fully tested without operating and destroying the device, and as a result failure-finding is impossible.

It may be infeasible to test a protective system if there is a significant chance of causing the multiple failure while carrying out the test. For example, testing a turbine overspeed trip by deliberately defeating the normal control system, or a tank high level trip by overfilling the tank may result in an unacceptable risk of the multiple failure occurring during the test. In this case it may be possible to bring the risk to within a tolerable level by careful wording of the task or by employing additional protection during the test. "Additional protection" does not need to be additional equipment: it may be that a second technician could monitor conditions and be ready to shut down the system if required. Although it can be good practice for a failure-finding task to replicate the abnormal conditions as closely as possible, it is essential to ask the following question before selecting the proposed failure-finding task.

"If the protective system fails to operate correctly during this test, is there a risk that the test will result in the multiple failure occurring?"

Even if it is possible to test a protective system, the required failure-finding interval could be impractical for two reasons: it may be too long, or it may be too short.

Long failure-finding intervals are common if the protective system is very reliable, demands on it are infrequent, and the consequences of failure are insignificant.

> A review group analysing a section of a chemical plant needs to set the failure-finding interval for a motor overload trip. The motor drives a water pump. The best estimates available to the group show that the mean time between failures of this type of trip is at least 200 years and demands are likely to occur no more than once every 20 years on average. The cost of checking the trip would be $30. If the trip failed to operate when required, the motor would burn out, but its replacement cost is no more than $250.
>
> The group determines that the trip's optimum failure-finding interval is 31 years.

When a properly calculated failure-finding interval is longer than the probable life of the equipment that it protects, the message is simple: be sure that it works today, then leave it.

Short failure-finding intervals are more challenging, and whether a specific interval is practical depends on the details of the system under analysis.

Incorporating failure-finding tasks in equipment start up or shut down procedures often provides the best opportunity for high frequency checks.

> Engine start up checks
>
> …
>
> After applying power, but before starting the engine, check that the following lamps are illuminated on the control panel: battery charging alarm; low pressure oil warning alarm;…

If a failure-finding task is required too frequently to be practical, it can mean that the design of the system is no longer able to deliver an acceptable level of risk. The redesign options discussed in the previous section should be considered.

3.9 Important Note

By now it should be obvious that the subject matter of this book is a very small part of a complete maintenance strategy. We have limited ourselves to hidden failures and mostly ignored those that are evident. The remainder of the book further assumes that the hidden failure cannot be managed by preventive or predictive tasks, so failure-finding is the only remaining option.

Remember that failure-finding allows the protective system to spend time in a failed state, unable to provide protection. For that reason it is important to consider options that *prevent* the failure before considering failure-finding. See other texts such as John Moubray's book (Moubray, 1997) for further information on failure management through fixed interval replacement, overhaul and condition monitoring.

3.10 Key Points and Review

Failure-finding is a task that checks whether a protective system is in a failed state. The protective system is allowed to run to failure, but its function is checked at fixed intervals to determine whether it has failed.

Failure-finding is not the only maintenance policy that can be used to manage hidden failures.

Because the protective device is allowed to run to failure, there is always a finite chance that it is in a failed state when a demand occurs on it. If failure-finding is chosen as a maintenance policy, there is a finite chance that a multiple failure will occur.

In general it is possible to manage the chance of a multiple failure by increasing the frequency of a failure-finding task, decreasing the demand rate on the protective device or both.

If condition monitoring, fixed-interval replacement or overhaul is technically feasible, it may reduce the risk of a multiple failure below the level that can be practically achieved through failure-finding.

If no maintenance policy can achieve a tolerable level of risk, the system may need to be redesigned to improve the availability of the protective system, reduce the demand rate on it, or to make the hidden function evident.