



*An Online CPD Course
brought to you by
CEDengineering.ca*

Mission Based Reliability: Turning Short-Term Survival into Long-Term Reliability

Course No: B04-006
Credit: 4 PDH

Daniel T. Daley, P.E. Emeritus



Continuing Education and Development, Inc.

P: (877) 322-5800
info@cedengineering.ca

Mission Based Reliability: Turning Short-Term Survival into Long-Term Reliability

Background

Ask yourself: What is the most important issue involved in taking a trip to Europe?

To me the most important issue is the survival of the aircraft for the flight to Europe and the flight home. Despite the importance of survival for these two “missions”, few travelers give this issue more than a passing thought. Why is that?

International travelers no longer worry about meeting their fate in the cold waters of the North Atlantic or during some other long trans-oceanic passage because the aircraft making those hundreds of trips each day are highly reliable. While there are some instances of incomplete transits, they are few and when they do happen, it is not the result of a failure of an established kind of aircraft using an established maintenance program. Failures result from circumstances that are well outside of the norm.

Generally speaking, the processes used to design and manufacture modern aircraft make certain that the inherent reliability of each aircraft is very high. In addition, the processes and programs used to maintain the reliability of the aircraft throughout its life are highly structured and highly disciplined in ways that ensure desired results are achieved. The maintenance management tools being used are fully aware of the useful life of critical components, the point at which each component is within its lifecycle and the manner in which each component will fail when it is nearing the end of its lifecycle.

Now for a moment switch your thinking from a situation in which reliability is nearly certain to a situation that is at the opposite end of the spectrum from the aircraft making long trans-oceanic trips. Think about a situation in which survival of the asset is not certain, in fact it is actually far from being certain. What would lead to this situation and are the issues leading to this situation the complete opposite of the characteristics described above for aircraft and the air travel industry?

The answer to that question is probably “No”. Even those assets that seem to have chronic reliability problems are likely to have been designed and constructed using reasonable processes and procedures and their maintenance programs are likely to be performing many of the tasks needed to keep the vast majority of the components healthy.

Many assets have thousands and thousands of components and the failure of even a small fraction of those components can make the asset seem to be a real reliability headache.

When dealing with poor reliability performance as described above, people often ask if the situation being experienced can best be described as a “million bee stings” or are the problems more akin to “a few snake bites”.

The real answer is probably, neither. The number and frequency of failures may seem as though something different is always failing, but most often it is the same few things that are failing repeatedly and unexpectedly. On the other hand, if it was only one or two items that were experiencing chronic failures, it is likely that the information and resources needed to solve the problem would have been found in the past and the problems would have been permanently solved before they could adversely affect the business.

So the correct answer is that you are probably not dealing with two problems and you are also not dealing with two hundred problems. In fact, in most instances, what seems to be a situation with an insurmountable number of chronic problems are those in which fewer than two- dozen different problems exist. In this discussion, we will refer to those two dozen items as “bad actors”.

But if an airline was experiencing a situation in which two dozen critical issues were repeatedly surfacing in a manner that ultimately led to breakdowns, two things would be true. First, you as an air traveler would have begun to think about the importance that survival plays in your trips and you would be reviewing itineraries for the kind of equipment being used during each segment of your trips. Second, the affected airline would be grounding the affected aircraft or would soon be out of business.

Introduction

Let’s begin by assuming that we are responsible for dealing with a situation in which an asset is experiencing repeated failures being caused by two-dozen different sources. While this situation may seem like a nightmare, it is one that is not uncommon. Many if not most new assets go through a process of commissioning and initial start-up in which a significant number of “bugs” must be worked out. Before the “start-up” is considered complete most of those bugs must be laid to rest and the asset operation must be in a stable condition. Individuals assigned to work on start-up teams are familiar with this kind of issue and they frequently find the hectic lifestyle fits their desire for excitement and need for adrenalin.

While most assets reach “stable” operations before being viewed as a truly finished product, there are instances in which either some of the “bugs” remain unsolved, some of the “bugs” were not recognized during initial operations or other “bugs” came to life sometime after sustained operation was achieved. In other words, there are instances in which two-dozen chronic bad actors can continue to be present in an asset after it is believed to have become a part of the fleet of stable income-producing assets.

In these situations, the operating and maintenance staff might come to believe that somehow they must have died and landed in hell. Something is always going wrong. The asset seems to

be always shutting down and asset managers are always receiving “helpful” phone calls from their superiors and home office asking what plans are in place and how soon the production stream will become steady and the asset will begin behaving as was expected when the investment of funds and other resources was made. Those kinds of questions are believed to help improve one's focus on the issues that exist.

The more important question is: Where is the foundation upon which we can build a program that will end this hemorrhaging of cash and other resources?

The answer to that question comes from one of the least expected places: The solution will be built upon the current reliability of the two dozen assets that have been causing all the problems. Even the two dozen components that have been causing all the problems due to short and unstable life-spans have a specific Mean Time Between Failure (MTBF). And their current MTBF can be used as a basis upon which survival and stability can be based.

There is an old saying: “It may not be much, but it’s not nothing”. That saying is true of a MTBF in spite of the fact that is much shorter than one would hope.

One of the most basic equations used in performing reliability analysis is the following equation:

$$R = e^{-t/MTBF}$$

Where R is the apparent reliability of a device, t is the interval between inspections or “health verifications” and MTBF is the Mean Time Between Failure of the component in question.

Let’s re-arrange the equation to produce one that better suits our current purposes:

$$t = \ln(R) \times MTBF$$

Or the required interval of a health verification is the MTBF times the natural log of the desired reliability.

For instances, assume we wished to increase the reliability of a component to 90% (or reduce the likelihood of failure to 10%), we could do so by performing health checks at an interval of:

$$t = \ln(.9) \times MTBF$$

Since $\ln(.9)$ is about $-.10$, we could achieve the desired apparent reliability by performing health checks at an interval of approximately 10% of the MTBF.

For instance, if the current MTBF is approximately 3000-hours, we could improve the apparent reliability to 90% by performing health checks on a nominal two-week interval.

While the above discussion provides the reader with the basic concept upon which much of the following methodology will be constructed, there are several other elements that will be applied in a process that will both increase the likelihood of near-term survival and expand the near-term survival intervals into long-term reliability.

Possibility the most important characteristic that is enabled using this approach is that of “stability”. Said another way, if it is possible to simply depend upon survival for the foreseeable future, it is also possible to do some of the other things that are impossible when the frenzy of repeated failures exists.

It will be easier for the reader of this document to embrace the efficacy of the concepts being presented if he or she has a mind-set of how those concepts will be applied within the system or process that depends upon them. For purposes of making the concepts more understandable, the reader should mentally place himself or herself in the following situation.

Begin by thinking about being in a situation in which you are the reliability engineer and are ultimately responsible for improving the reliability, availability and Total Cost of Ownership (TCO) of an asset or fleet of assets. Also assume that one or more of the assets is failing so frequently that it is difficult to sustain a focus on any other opportunities that are available. In fact, it is necessary to come to work early, work late and work some portion of almost every weekend just to keep your head above water.

In your role, you are currently unable to do the kinds of things that many other reliability engineers are doing. You are not extending the useful lives of many components from one-year to two-years, then from two-years to four years, and so on. You are simply struggling to achieve day-to-day and week-to-week survival. No sooner do you begin to see some “daylight” that another round of failures begins and existence again becomes a struggle.

Let’s assume that in response to that way of life, you have decided to try to make improvements using Mission Based Reliability (MBR) by focusing on those things required to achieve near-term survival.

The first step of achieving improvements using this technique is to form and train the team whose job it is to identify and take all the steps needed to ensure the survival of the asset targeted for improvement.

For purposes of this discussion we will initially assume that the optimum mission interval is two-weeks so the meetings of the team and all activities they cause to occur will happen on a two-week interval.

The bi-weekly interval will begin every other Wednesday. Before the Wednesdays on which the meetings are held, the team members will gather all the things that should be addressed to ensure the asset will survive without failure for the period between the next PM and the following PM, two weeks later.

The first thing that must be done is the complete list of simplified health-checks for the two-dozen items on the bad-actor list. If any of those items are found in a deteriorated state or on their way to failing, the decision must be made if the component in question can survive until the next bi-weekly PM or if it must be changed immediately. In addition to the health checks for the bad actors, the complete list of normal predictive and preventive maintenance must be done.

Finally, the team must discuss any warnings or complaints they have received from the members of their operating and maintenance teams. The reliability engineer must identify the activities needed to investigate those complaints.

The Thursday following the Wednesday meeting is spent preparing for the PM activities that will occur on Friday. On the Friday after the bi-weekly Wednesday meeting, the entire list of periodic maintenance activities will be completed by a multi-person team with the number of participants determined by the time required to perform those tasks and the number of work-paths that can be worked simultaneously. A simplified form of Critical Path Planning may be used to plan and schedule this work. Again the scope during the Friday PM will consist of:

1. Streamlined health checks for the two-dozen bad actors.
2. Normal predictive and preventive maintenance.
3. Follow-up on warnings resulting from unusual noise or other symptoms reported by operators and maintenance personnel.

Following the Friday PM, the asset is expected to operate in an uninterrupted fashion through the next two weeks. On the Wednesday two-weeks after the last meeting of the "survival team", the entire cycle will start once again with another survival team meeting setting the stage for Thursday preparations and Friday PM.

While it is intended that the asset operate without any form of interruption, there are likely to be instances during which the heightened sensitivity and surveillance of operators and crafts will identify situations deserving of immediate response. In those cases, decisions will be made if the situation can be trusted to survive until the next planning cycle or if it must be addressed now.

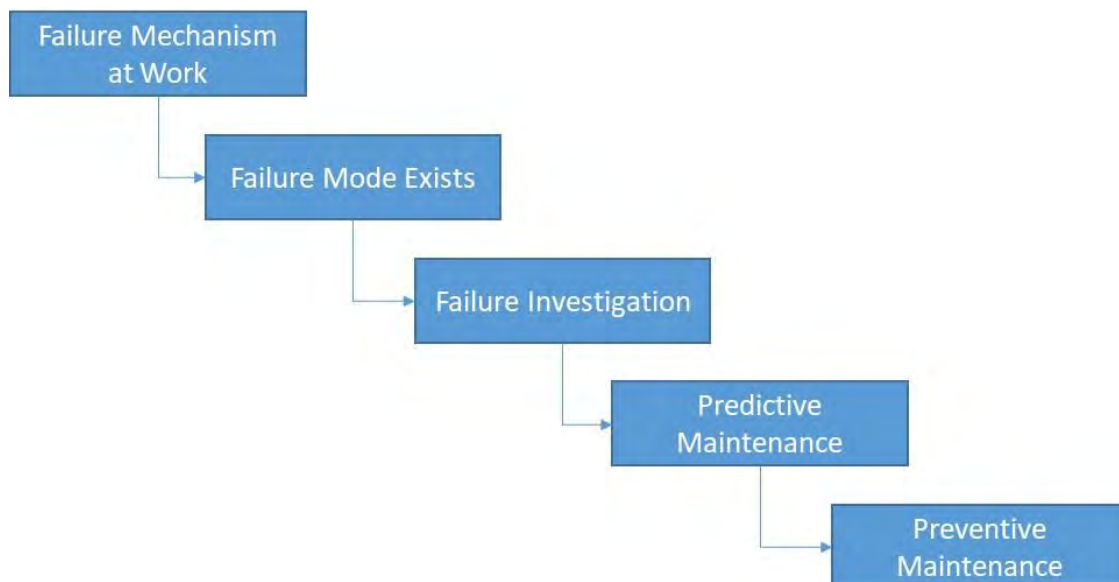
While this process may seem to require far more activities than are currently occurring, the events are structured in a way that prevents them from creating interruptions to everything else that should be happening, thus allowing normal progress to occur.

Critical Connections in a Normal Situation

In a normal situation, the steps leading from the start of deterioration to the replacement of a component follow the pattern described below:

1. The physical circumstances that create the presence of a **Failure Mechanism** are introduced into the asset or system. There are four general kinds of Failure Mechanisms: Corrosion, Erosion, Fatigue and Overload. When the conditions are right for the existence of a Failure Mechanism, deterioration begins to occur.
2. Deterioration of some specific element of a component occurs. Ultimately the deterioration reaches the point that the component can no longer perform its intended function. At this point, the **Failure Mode** exists and the component fails. In the normal life-cycle, the deterioration and failure occurs over a period and at a frequency that fulfills the reliability promised by the manufacturer.

3. In the typical cycle of creating some form of prevention, the Failure Mode is reviewed. The condition of the component upon failure provides the clues needed to identify the exact **Failure Mechanism**. If failures of the kind being investigated have been present for some time, it is possible to identify the Mean Time Between Failures of this component, its shortest life-span and its affects. With that information in hand, it is possible to identify the specific form of **Predictive Maintenance** needed to identify the presence of deterioration, decide when Predictive Maintenance should be started (e.g. shortly before the shortest life span) and decide when the predictive maintenance should trigger replacement of the component.
4. While performing Predictive Maintenance leaves the component “Good as Old” (meaning no improvement is made), you ultimately wish to eliminate all the deterioration and return the asset to “Good as New” condition. This is done by replacing the deteriorated component and doing so before a failure is allowed to occur. While it is intended that the Predictive Maintenance trigger the timing of **Preventive Maintenance** (or replacement of the deteriorated component), there are instances when the cost of the effect caused by the failure is so great that you wish to avoid the failure in all cases. In this situation you may wish to trigger the Preventive Maintenance by additional factors than just the Predictive Maintenance. For instance, if the Predictive Maintenance has not triggered replacement by a specific point in time, you may wish to trigger the Preventive Maintenance to occur at the MTBF or 90% of the MTBF if the effects of a failure are particularly distasteful.



Critical Connections in Mission Based Reliability

While the techniques used as a part of Mission Based Reliability are much the same as for normal forms of prevention, for the two-dozen bad actors being addressed by MBR, the lifecycle of the

components are not nearly as long and are causing far too many nuisance breakdowns. As a result, the Critical Connections as they are addressed in MBR are described as follows.

1. **Failure Modes** tend to occur earlier and more frequently in the life-cycle of certain components.
2. As in the normal case, it is necessary to identify the MTBF despite the fact it is short and likely unstable. It is also necessary to develop an understanding of the P-F interval and, if not the **Failure Mechanism**, at least the signs of deterioration and symptom of impending failure.
3. Rather than typical **Predictive Maintenance**, it is intended that streamlined forms of inspection or health-checks be applied. Generally speaking, MBR inspections are applied more than six times as often as normal Predictive Maintenance so they need to be designed to be much more efficient while still being effective.
4. Unlike more sophisticated forms of Predictive Maintenance, MBR inspections and health checks are looking for more evident signs of impending failure. The equivalent of the P-F interval between these signs and the failure they are signaling may be shorter. As a result, when an impending failure is signaled, it is important to make the decision if the **Preventive Maintenance** must happen NOW or if it can wait the short time until the next opportunity for replacement.

Basic Tools of Mission Based Reliability

What are the key concepts upon which Mission Based Reliability is based? What things must be included in the training?

The first concept is the one described above. It is possible to improve the apparent reliability of a component by leveraging the current MTBF to achieve a much higher level of reliability than would be provided by the MTBF alone.

Let's assume that we will be satisfied with 90% reliability during the two week intervals between health checks.

We will begin by identifying the MTBF of the two-dozen components that are experiencing problems related to shortened life. Once the MTBF for those two-dozen components has been determined, we continue by calculating the required interval for the health checks for each of the two-dozen bad actors using the following simplified equation:

$$t = 10\% \times \text{MTBF}$$

This inspection interval will deliver an apparent reliability of 90% or more.

For simplicity we will assume that all values for t will be near the two week interval or longer.

Once we have determined the inspection or health check interval for all of the bad actors, we create a set of simplified tasks that are to be used to determine that each component is still functioning and there are no obvious signs of deterioration or impending end of life.


It is important for the members of the survival team to understand how this method works to improve the apparent reliability of any component. If you perform the simplified "health check" on a frequent interval, the component can be replaced either when it is found to be failing or at the next opportunity but still before it fails. The frequent inspections cause the number of items that are allowed to fail without replacement during the short mission intervals to be quite small or zero.

The survival team will be made up of the Operations Supervisor, the Maintenance Supervisor and the Reliability Engineer. This small critical mass of individuals must understand enough about the simplified statistics and physics-of-failure to take whatever actions are needed to ensure near-term survival of the asset.


The next tool the survival team must understand is the concept portrayed by the P-F interval. In P-F, the P stands for "potential" and the F stands for "failure". So the P-F interval portrayed in a P-F curve represents the path between when an event or condition that dramatically increases the potential for a failure occurs and the occurrence of the failure.

Few if any P-F curves present actual data precisely representing the likelihood of failure for each point on the curve at each specific point in time. But a P-F curve can be used to provide a meaningful representation of the change of the rate at which the likelihood of survival decreases leading up to the failure.

Say, for instance, you are dealing with conditions that set the stage for external corrosion of a structural member (for example the introduction of an electrolyte in the presence of dissimilar metals). In this case, degradation occurs slowly and the portrayal of a P-F curve will be as follows.



Now, let's assume that corrosion is taking place inside of a circulating system and that there is a filter meant to capture all the corrosion products before they reach the small clearances of engine bearings. Assume that once the filter is completely loaded, the increased pressure differential across the filter will cause the filter to collapse and all the corrosion products it contains will be released into the fluid stream and reach the bearings. The effects of this kind of failure are almost instantaneous and the P-F curve would be drawn as follows.



This is the kind of knowledge your survival team should understand. They should clearly understand that some of the two-dozen bad-actors will fail in a gradual manner while for others the failure pattern will be nearly instantaneous.

As the team does its work of taking the steps needed to ensure survival, they need to have at least a basic understanding of the physics-of-failure for different kinds of failure modes and they should make decisions and take steps that recognize those differences.

In addition to the intuitive sense of survival and failure described above, it would be useful for the survival team to understand several other concepts.

While the concept of “bootstrapping” may seem different from what is being accomplished by using MBR, there is some applicability of the concept. In general, bootstrapping is an approach by which a person is able to elevate himself or herself to some higher position, level or capacity literally by leveraging the small skills or capabilities with which they start.

The concepts we are applying in this process is, in fact, a form of bootstrapping. We are attempting to use both a limited amount of knowledge and a small but still present degree of reliability to achieve much greater performance.

The following characteristics form part of the basis upon which this bootstrapping process can be achieved:

- We know that increased surveillance will result in improved "apparent reliability"
- While not exact, the concept portrayed by a P-F curve and its shape can help us understand how quickly failure will occur after the potential for failure is present
- Simplified inspections can be used to provide an adequate warning of incipient failure resulting from short to moderate interval failure modes

Other useful tools include:

- Close physical proximity can be used to our advantage when assessing the condition of a component
- Hyper-sensitivity of operators and maintenance personnel to unusual noises, operating characteristics or symptoms of deterioration can help identify newly developing problems
- Changing patterns and relationships representing anomalies can be used to identify and further understand newly developing problems

Individuals, who possess sufficient intelligence and sensitivity to successfully operate a complex asset between eight and twelve hours per day, are also able to sense some things that those, who either do not share that sensitivity or are not exposed to the operation of an asset for nearly as much time, are unable to identify. When leveraging available assets, it is important to take full advantage of the sources of information closest to where the failures exist and where they exhibit their symptoms.

In making use of the concepts and resources upon which we will leverage our efforts, the following definitions will be helpful:

- Mission Based Reliability or MBR - For the purposes of this course, Mission Based Reliability is a process that depends on repeatedly taking the steps needed to ensure the survival of an asset for the relatively short interval needed to complete a defined mission interval.
- Mean Time Between Failure or MTBF - The Mean Time Between Failure is the average of the time between when a new component falling into a specific group is installed and when it fails. While the MTBF can be used in calculating the life-spans of an entire population of equipment types (like all pumps), it is most useful when focusing on a component type having the same manufacturer, model number, application and severity of usage.
- Inherent Reliability or Actual Reliability - The inherent or actual reliability of a component is the measure that represents the useful life and/or failure rate that is achievable without any assistance or human intervention that helps hide or prevent failures. The inherent reliability of a component may best be found by the life span it will experience when it is simply installed, then left to run to failure.
- Apparent Reliability - The apparent reliability of a component is that level of performance that appears to be present when steps are taken that helps in preventing failures that might have otherwise occurred. For instance, frequent inspections will identify instances when the condition of a component is failing. When it is found that the component is failing, it can be replaced before failure so the apparent reliability seems much higher.
- Pre-Departure Testing - Pre-departure testing is a set of inspections or tests to which an asset is exposed immediately before it begins to perform a mission or before a specific mission interval. The Pre-Departure tests identify any critical components that may have experienced an event raising its potential for failure or having failed without discovery. When deteriorating components are found, they are evaluated to determine if they will survive the mission. If they are unlikely to survive the mission or mission interval, they are replaced.

- P-F Interval - The P-F interval is the somewhat inexact interval between when the potential for failure is present in a component and when failure actually occurs. Not all P-F intervals are the same, so to avoid failure, the P-F interval must be assessed to determine when action must be taken to avoid failure.
- Short Interval Failures - Short Interval failures are those that are likely to occur within a very short time after the potential for failure is present.
- Long Interval Failures - Long Interval failures are those for which there is a significant amount of time between the enabling event and the failure. For instance, if a situation that introduces an exposure to a dielectric (water) is the initial event that increases the potential for failure, the resulting failure will not be instantaneous. There is time to respond and still prevent the failure from occurring.
- Mission Interval - While the mission interval for an aircraft flight or a train route is clearly the time between departure and arrival, the mission interval for a stationary asset, like a plant or refinery, is a period that is created to both meet the needs of the inspection intervals and the rhythms of the organization where the asset is located.

Mathematics of MBR

Returning to the equations described above:

$$R = e^{-t/MTBF}$$

and

$$t = \ln(R) \times MTBF$$

It is useful to create a table that displays the trends that are produced when dealing with various MTBF levels and various desired levels of reliability.

		Inspection Interval (t-hours)					
		Desired Apparent Reliability					
		70%	80%	90%	95%	99%	99.90%
Current MTBF (hours)	In of R	0.3567	0.2231	0.1054	0.0513	0.0101	0.0010
	1000	357	223	105	51	10	1
	2000	713	446	211	103	20	2
	3000	1070	669	316	154	30	3
	4000	1427	892	422	205	40	4
	5000	1784	1116	527	257	51	5
	6000	2140	1339	632	308	61	6
	7000	2497	1562	738	359	71	7
	8000	2854	1785	843	410	81	8
	9000	3210	2008	949	462	91	9
	10000	3567	2231	1054	513	101	10
	20000	7134	4462	2108	1026	202	20
	30000	10701	6693	3162	1539	303	30

When reviewing the table, several features should be noted:

- As increasingly higher levels of reliability are desired, the required inspection or health-check intervals become smaller to the point the interval is impractical when using manual means of inspection.
- Very high levels of reliability are impossible independent of inspection interval for items with very small MTBF. This highlights the importance of the inherent reliability present in the components being used.
- When health verification times become too short, the only methods that can perform at the required frequency is by using electronic continuous monitoring.
- Note that one year of round-the-clock operation is 8760-hours, so 9000-hours is a MTBF of slightly more than one-year. 20,000 hours is about two and one-quarter year and 30,000 hours is about three and one-half year. It is useful to have a sense of the MTBF plateau needed to survive using inspections conducted on a normal interval.
- Note that 90% reliability can be achieved using health checks conducted at an interval slightly greater than 10% of the MTBF. Even items with a MTBF of approximately three and one half years, health verifications must be performed on an interval somewhat more frequently than semi-annual to provide a reliability of greater than 90% (or a likelihood of failure less than 10%). Said another way, even those parts viewed as being “reliable” need to be included in quarterly or semi-annual health checks.

If we were to use this chart to help identify a useful “rule-of-thumb” for health-check intervals, we would focus on the inspection intervals required for 90% reliability for items having a MTBF of 3000-hrs or 4000-hrs. A two week interval is 336-hours, so the two week interval between health verifications may be an appropriate interval for many components with those MTBF levels.

There are a few important observations that can be made concerning frequent inspection intervals:

1. In addition to finding and preventing failures, increased surveillance is likely to improve the understanding of the Potential-to-Failure relationship by increasing your understanding of the period between when potential-causing events occur and when the associated deterioration is first recognized. An increased understanding of the circumstances that causes deterioration and failure can lead to improved prevention by avoiding those situations.
2. Increased surveillance can also provide insights into when “life-extending” tasks are needed. For instance, when overheated, contaminated or discolored lubricants are quickly identified, it is possible to replace the fluids before harm is done to permanent components.
3. A well-designed set of inspections of the two-dozen bad-actors will bring knowledgeable personnel in close contact with many areas of the asset. As a result, inspectors can see if other forms of deterioration are occurring in other areas of the asset.

Designing a Mission (or Selecting the Mission Interval)

A mission consisting of an aircraft flight over an ocean designs itself. The mission starts with take-off and ends with arrival. Some trans-oceanic flights start at coastal cities while others start at cities far from the coast, so there can be thousands of miles over land before the segment over water. The same is true of the destination. Some destinations are at cities close to the coast while others are at locations far removed from the coast.

If you are a skittish traveler, you may find some comfort in knowing you are currently travelling over land, this degree of comfort is not shared by the airline. If a flight has to be curtailed because of a problem that occurs over land, the fact is that the problem just as easily could have occurred over water. For the airline and the aircraft manufacturer, the failure has to be equally worrisome wherever it occurs.

This same kind of sensitivity to failure needs to be applied when designing a mission interval for an asset that is stationary or for which a failure will not result in an affect as dramatic as a crash into an ocean. You need to design a mission interval during which the statistical likelihood of survival remains high enough to meet the requirements that have been set.

If the desired mission reliability is 90%, then the KPIs used to measure success or failure of the program should indicate that fewer than one-in-ten mission intervals have resulted in failures during the mission. In other words, when utilizing two-week mission intervals, you should experience only slightly greater than two mission interruptions per year.

Compared to a situation in which two dozen bad actors are being allowed to trigger breakdowns on their own statistical failure schedules, the situation described above is a dramatic improvement.

While the above paragraphs provide the reader with a sense of the importance of properly designing the mission interval they do not answer the question: How do you design the mission interval?

Analyzing the failure pattern of the two-dozen bad actors will provide you with the MTBF of all those components. Assuming that you wish to improve the apparent reliability of those items to 90%, you can multiply the MTBF for each item by .10 to determine the interval at which you should perform health-checks for each of the two-dozen or so bad actors.

With that information in hand, the next step is to create a rational program for performing both the health-checks and for performing the preventive maintenance tasks (changing the components) that will be triggered by the health-checks. (Keep in mind that you will want to replace each failing item within the remaining survival time before failure. These replacements will preferably occur during opportunities that will not cause additional down-time.)

When taking this step, you will quickly realize that a factor bearing on the mission interval that may be more important than the mathematics and the component analysis is the rhythm with which your organization functions. While it is possible to design requirements for performing inspections and maintenance on almost any schedule, it is unlikely that the work will be done in a timely way unless the schedule fits the rhythm with which your organization functions.

The two-week mission interval introduced earlier is one that fits many organizations, but not all. The logic for using the Wednesday-Thursday- Friday pattern for reviewing current survival requirements, planning and scheduling the work then performing the work is an approach fits many organizations. Many of the decisions are made by staff personnel who typically work on day shifts. Mondays and Tuesdays of the weeks in which “survival meetings” are held are frequently busier than other days later in the week but do provide an opportunity to catch-up with issues that occurred over the weekend and an opportunity to determine if new problems have appeared.

If the MTBF of many of your two-dozen bad actors are less than 3000-hours, your mission interval may need to be less than two-weeks. If the MTBF for the majority of the two-dozen bad actors is much greater than 3000-hours the mission interval may be longer than two-weeks.

Some organizations may work on schedules in which the typical seven day weekly pattern does not exist. Or they may be organizations that do not predominantly staff day-light shifts with engineering staff and supervisors who will comprise the “survival team”. In those instances, the entire pattern of the survival meetings, health-checks and PM can be set to whatever rhythm best fits the human organization.

The following are a few of the characteristics that were considered when designing the two-week mission interval. If they do not exist in your case, you will need to substitute the characteristics that are important to you:

1. The two-week interval is based on the working rhythms associated with a typical 5-day work week.
2. It is based on key staff members of the “survival team” all working together on daylight shift on Wednesday, Thursday and Friday.
3. It assumes that while the members of the “survival team” already have full-time jobs, they will have time for participating in well-structured survival team meeting on an every-other-week basis.
4. It assumes that effective predictive and preventive maintenance programs already exist for components and systems other than the two-dozen bad actors and that they can be integrated with the new work-load.
5. It is also assumed that a small number of the health checks will need to be completed more frequently than every-other week and that these items can be accomplished by the existing personnel on a weekly or more frequent interval.

Pre-Departure Tests, Health-Checks, Simplified Inspections

Earlier, the concept of Pre-Departure Tests, Health-Checks and Simplified Inspections have been mentioned. For purposes of this discussion all of those items are much the same.

The general concept is that the health of a component can be determined quite quickly and accurately if an individual who is intent on performing the evaluation is placed in close proximity with the component and has a clear description of the factors upon which the evaluation is made.

Clearly electronic black boxes show very few external signs of deterioration. All that can be determined is if they are providing the desired output based on a specific input. Either they are operating or they are not. There may be instances in which the response is erratic or inconsistent, but even that response should be viewed as a failure.

On the other hand, many mechanical items either have signs of physical deterioration or external symptoms that indicate internal deterioration. For example, if a pump vibrates more than normal, runs hotter than normal or fails to put up the expected “shut-off” head (discharge pressure with discharge valve closed) it is possible to say that a failure has occurred or is imminent.

Sensors indicate failure when the output does not reflect the normal input.

Wire-ropes have deteriorated to the point of replacement when a specific number of broken strands are present in each specific length.

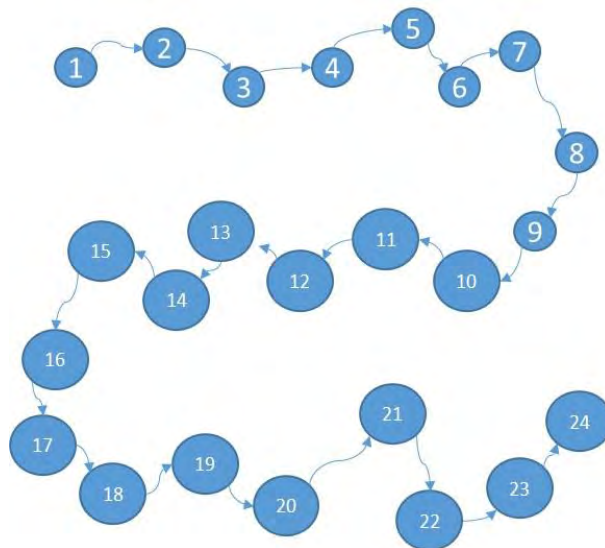
While it requires some thought for each of the two-dozen bad actors, it is possible to identify a specific form of simple inspection or test for almost every device. If the requirements are clearly described, it should be unnecessary to spend more than a minute or so making the determination if an item is “healthy” or if it should be replaced.

Organizing Health-Checks

Health-checks, simplified inspections, and other forms of predictive and preventive maintenance that must be completed in brief periods can be significantly improved by additional structuring and organizing. It is helpful to create a sequence for tasks that allows the person performing the inspections to go from one task to the next closest task in sequential order. Maintaining the focus on the sequence of tasks is helped by marking the tasks with sequential numbers shown on decals mounted close to the point where the inspection is conducted.

In some instances, specific tasks may be part of more than a single sequence and in that case, there might need to be several sequence numbers associated with certain tasks. In this case, the decals affixed to tasks for different sequences may need to have different colors.

For tasks that are particularly complex, it might be helpful to affix a decal describing details close to the place where the task is accomplished. An example might be the frayed wire count and length over which counts are made that are used to condemn worn wire ropes.



When the work is organized in this fashion, it is possible to create a simple checklist for reporting components that are acceptable and those that should be replaced and the seriousness of the situation.

MBR - Health Check Record			Date: __/__/__
Item Number	Health-Check Description	Condition Acceptable? (Y/N)	Level of Urgency (Now/Next Opportunity)
1	Visual	Y	
2	Function Test	Y	
3	Visual	N	Now
4			
5			
6			
7			
8			

Decisions Based on the P-F Interval

While the actual amount of time between potential-for-failure and failure are not shown by P-F curves, it is helpful for individuals making replacement decisions to know if the P-F interval for specific components is short or relatively long.

For instance, if a black box or electronic component shows signs of unstable or incorrect output based on an input, it should be expected that the component is likely to fail without further warning.

A pump that has a seeping seal is likely to continue to deteriorate on a gradual basis until a major leak develops. In this case, the replacement should be made at the next opportunity that will not result in additional downtime.

If a pump fails to put up the desired shut-off head, but is capable of performing the currently needed pumping rate, the pump should be replaced and overhauled at the next opportunity that is already part of an extended outage.

Associating a realistic P-F curve with each of the two-dozen bad actors will help both the survival team and the individuals performing the inspections to make better decisions concerning how quickly a failing component must be replaced when specific signs of deterioration are evident.

Key Elements on Which to Focus

When performing one of the several forms of reliability analysis, analysts soon learn to distinguish between critical and non-critical components and between dynamic and static components.

Each asset typically performs between a few and a dozen functions. Critical components are those that are needed to perform one of the functions being performed by the asset. Non-critical components are those that can fail without resulting in loss of one of the functions. As a result it is important to focus on deterioration of critical components.

Dynamic components are those that move in some manner or create a change between incoming and out-going signals or streams. Static components are those without moving elements or those that do not produce a change between incoming and out-going streams or signals. Generally speaking, much more attention must be paid to dynamic components than static components.

If you were to identify the characteristics of the two-dozen bad actors, it is likely that all of them are both critical and dynamic components.

When monitoring components to identify surfacing issues that may turn into failures, it is most important to spend the majority of time focusing on components that are both critical and dynamic. Those are the items for which a failure will most likely prevent survival of the asset.

These items will include:

- Pumps
- Compressors
- Motors
- Turbines
- Valves
- Sensors
- Actuators
- Controllers
- Bearings
- Seals
- Switch gear
- Transformers
- Etc.

Recognizing and Integrating Emerging Issues

For a moment, focus on the components that are not among the two-dozen bad actors. As suggested in the last section, apart from performing the normal forms of PM on these items we will spend some of our energy focusing on any signs of incipient failure that will not be addressed by our current programs of predictive and preventive maintenance.

While experience with an asset typically leads to a situation in which the owner's personnel are aware of the normal patterns of lifecycle and failure and have introduced programs that accommodate those patterns. Unfortunately, it is not uncommon for things to change over time. For instance:

- Failure modes that have been eclipsed by more dominant failure modes may begin to become apparent when the current Dominant Failure Mode is solved. For instance, if a component is experiencing both corrosion and fatigue and has always failed due to corrosion before the number of fatigue cycles have caused a problem, if the frequency or stress level of fatigue increases, the failure mode resulting from fatigue may begin to predominate.
- Assume new or different operators are assigned to an asset. Assume the manner in which they operate the asset places greater stress on the asset. Then it is possible for a new failure mechanism to be introduced.
- Assume the supplier of replacement parts is changed or assumed that the current supplier changes his manufacturing or quality control processes, a new failure mechanism could be introduced or a former failure mechanism may be intensified causing more frequent failures.
- Assume that the individual providing the vigilance needed to identify deterioration become lax, then it is possible that deterioration will go undetected and failures will be allowed to occur.

In these instances, we need to remind ourselves that the sole objective of the survival team is achieving survival, no matter what it requires.

Let's think about the physics-of-failure for these items. If they are being addressed using the conventional predictive and preventive maintenance program, we are counting on that program to identify deterioration and trigger replacement before failure occurs.

For instance, if the MTBF of a component is three years, and the earliest failure in the past has occurred at the two and one-quarter year point, we may have decided to begin the quarterly predictive maintenance at the two year point. We expect that the deterioration rate is slow enough that the P-F interval is much longer than three months so when we find deterioration, we have more than three months to respond.

Now let's that think about the kinds of changes that may make that thinking and the resulting strategy ineffective.

- Say, the quality or durability of the replacement components deteriorates, just a little, and the earliest point of failure moves to slightly less than two years. Now there will be no predictive maintenance to trigger the needed replacement.
- Say the severity of stressors increase, just a little. Then the P-F interval will become shorter. If we continue to believe we have a three month period to make replacements after the deterioration is discovered, the current delay in response will allow the component to fail before replacement. In this instance, if we have set up our on-hand inventory in a way that orders a replacement only after deterioration has been identified, the failure will occur with no replacement part on hand.

While the two-dozen bad actors are the primary focus of the survival team, all the other critical and dynamic components in the asset should also be part of their concern.

While focusing on the two-dozen bad actors will provide members of the survival team with a much clearer idea of how components fail and how the process of prevention works, it is fully expected that those learnings will ultimately carry over into a level of sensitivity and vigilance for all the components.

Follow Up on MBR Inspections

If a simplified inspection of one of the two-dozen bad actors recommends that the component be replaced, what should you do?

If your answer is that you should replace the component, you are probably right. On the other hand, if it is early in the program, the simplified inspections being performed may be either too conservative or not conservative enough.

Early during the implementation of the MBR program, it will be necessary to calibrate the simplified inspections relative to the actual conditions being found and being recommending for replacement. If you find that the components being recommended for replacement are still good, you may wish to adjust the level of sensitivity of the inspector. Obviously, the component being replaced will always have some amount of useful life remaining (or you would have found it in a failed condition). But if the component has only cosmetic deterioration, the incipient failure condition thought by the inspector to be incipient failure should be adjusted.


On the other hand if you simply make allowance for the conditions at which replacement is being recommended by ignoring the recommendations, the person performing the inspections and making recommendations will adjust his or her performance by providing reduced vigilance and fewer recommendations. This situation will destroy the effectiveness of the MBR program by eliminating the trust that must be implicit in the process. The survival team must trust that inspectors and other operators and maintenance personnel will report conditions that seem to be on their way to failure and those individuals must trust the survival team to take the appropriate action when a recommendation has been made.

The Reliability Model or Anticipated Results being Produced

The reliability model for a component is a highly summarized and simplified mathematical description of the likelihood of failure of each component over time.

Reliability engineers are used to using statistical distributions, like a Normal distribution or a Weibull distribution, to describe the likelihood of failure over time. For instance, if the life of a component is best described by a normal distribution, and if a population of components are

allowed to run-to-failure, the number of failures at each point in time would resemble a typical bell curve.



It is useful to envision the reliability model that will be produced by the MBR process. When components fail with a very short MTBF, it seems unlikely that the pattern of failure (in a run-to-failure mode) would present a “mature” distribution as in the case of the bell curve above. In the case of MBR, the frequency of inspections and the timely replacement that is triggered is intended to produce a specific reliability or likelihood of survival. In the examples most commonly used in this discussion, we have used the inspection intervals needed to provide a reliability of 90% (or higher).

As a result, we can assume that the results expected from MBR would be a constant reliability of 90% or higher. The term “or higher” is being added because, rather than the results of a simple predictive maintenance – preventive maintenance regimen, MBR is being done in a highly interactive manner with near continuous intelligent and sensitive vigilance being applied to identify and respond to known and arising issues to ensure survival. The likely results are therefore expected to be higher than the calculated level.

Alternate Outcomes from MBR

Once the MBR process is started, it is fully expected that the simplified inspections will identify failing components among the two-dozen bad actors before they fail, thus providing an opportunity to replace them without experiencing a breakdown. In spite of that objective, there are a number of other outcomes that may result:

- Bad actors may continue to fail as they have in the past. In this case, the simplified inspections are not effective. In this case, it may be necessary to either re-design the simplified inspections or re-emphasize the way in which the inspection should be conducted.
- The simplified inspections find nothing, trigger no replacements, but some or all of the bad actors simply stop failing. Bad actors begin to behave like the components being maintained by the normal maintenance program. This effect is not unusual. When a light

is focused on a problem, it is not unusual for the increased level of attention alone to result in improvement. In this case, it is useful to determine if the increased vigilance has resulted in the desired level of reliability. If so, it will be useful to develop a method of more widely using this method while continue doing it as efficiently as possible.

- An increasing number of “opportunities” are identified and the list of components being treated as bad actors increases. While this may seem to make the process more cumbersome, it is actually an endorsement of the process and its benefits. Again, this result suggests the need to expand the program in the most effective and efficient manner possible.
- As part of the simplified rounds being conducted, the persons performing inspections identify other things that he thinks should be inspected. Like the item above, this is once again an endorsement of the process and, in that respect, is positive. On the other hand, if it is being suggested that the MBR program and personnel be used to provide oversight for conditions (like cleanliness) that should be observed by others as part of their duties, expanding the program to issues unrelated to survival is likely to dilute the focus on survival with bad affects.
- Some combination of all these things happens. It would not be surprising to see all of these issues at some time or other over the life of the program. If so, the issues will need to be managed by the survival team as they appear.

Structuring MBR for Long-Term Results

In the introductory discussion, we highlighted the effectiveness of a MBR-like approach when used to manage and ultimately solve a significant number of problems when dealing with the initial start-up of new assets. We also discussed the use of Mission Based Reliability to help ensure that equipment like planes and trains that continually perform missions are regularly able to finish their missions without a failure. A remaining question is: Will this approach be effective when applied to "created" mission intervals (e.g. those used in managing stationary assets like plants or mobile assets that are not isolated from sources of maintenance during their mission intervals).

Answering that question will depend on how well you are able to create the illusion or the belief and commitment to survival among the Survival Team and those they depend upon. If they are convinced that survival is critical to their own success and the success of both their organization and their company, then it will be possible to achieve a very high level of survival.

On the other hand, if they are allowed to believe that this is just another program or that survival is just another among a myriad of co-equal goals, it will be impossible to achieve the desired results. While this is most true at the beginning when the program is first being implemented, it remains important as long as the program is being applied.

When survival is treated with the same level of importance as other typical day-to-day issues, it will be prioritized along with them and on some days it will receive only luke-warm attention.

Breathing Room to Make Improvements Resulting in Less Dependency on Constant Vigilance

When the MBR process is successful, the results will provide some breathing room or an opportunity to make permanent improvements to components.

While it is possible to create a level of excitement and enhanced vigilance among your personnel for a reasonable period of time, if your assets are constructed in such a manner that they depend on this strategy for the long-term, it is likely that sooner-or-later, you will be unable to sustain the vigilance. As a result, it is important to use the opportunity provided by the initial success to make permanent improvements that allow the predictive and preventive maintenance needed to sustain reliability to be completed on a normal schedule.

Ultimately, it will be important to use the time created by the decrease in breakdowns to either provide forms of immunity to the dominant failure mechanisms or provide normal forms of prevention using predictive and preventive maintenance or both to place these items on a less demanding schedule of attention.

Improving the Degree of Robustness

While two-dozen items with relatively short MTBF is a large enough number to make an asset seem very unreliable, it is important to realize:

- Of the several thousand components of which an asset is constructed, only a few tenths of a percent are experiencing frequent failures.
- Even the two-dozen bad actors are performing their intended function, although not at an acceptable lifespan but largely because of a single dominant failure mode.
- Knowing how to provide the required functionality being provided by the two-dozen bad actors and knowing which their features that are currently able to survive and those that do not survive, provides a strong basis upon which improvements can be made.

The point being made here is that the same knowledge and understanding used to design and build the asset and to create an effective inspection and maintenance program for all the other components is precisely the knowledge needed to improve the reliability of the two-dozen bad actors. That information provides both the ability to repeatedly achieve short-term survival while, at the same time, providing the knowledge needed to make the kinds of changes that will significantly improve the MTBF and actual reliability of the two-dozen bad actors.

Said another way, it is necessary to improve the robustness of roughly two dozen components with respect to a single form of deterioration for each. If the asset can be kept operating while those bad actors continue to have short life-cycles, there should be time to focus on those issues and find solutions. Once a component has been improved, it will removed from the bad actor list and then maintained with normal interval predictive and preventive maintenance.

An Issue that is Frequently Confronted in MBR

While it is always important to compare the useful life of components and their MTBF to the useful life and failure rate promised by the manufacturer or supplier, it is particularly important to do so in the case of the two-dozen bad actors included in the MBR program. It is far more likely that these components are not fulfilling reliability promises and that the manufacturer or supplier is in some way liable for the costs of replacing them.

In addition to the expected useful life and the MTBF (or failure rate), manufacturers frequently provide the LTPD (or Lot Tolerance Percent Defects) expectations for their products. Based on a combination of the results of HALT testing, ISO 9000 procedures being applied during their manufacturing and HASS testing of finished products, the manufacturer should understand the percentage of “bad” products that should be reaching the customers. While the numbers of components you are using may not provide a large sample, your small sample is typically large enough to provide a statistical comparison with a small confidence level.

For instance, if you have used ten of a certain component and five have failed with a life-span far less than the manufacturer has promised, and the manufacturer has indicated that the LTPD should be 1%, then you have a right to suspect that the manufacturer has a manufacturing or quality control problem. Rather than 1% failures, your small sample is suggesting 50% failures. You should ask the manufacturer for the results of HASS tests for the lot or serial number interval containing the items you have received.

When the current situation is driving you to include a specific component in the group of bad actors being monitored by Mission Based Reliability, there is automatically a reason to include the component manufacturers in the root cause analysis process and in the corrective action process.

MBR Training

There are a number of skills that the individuals who will be involved in the MBR program will need to have. It is important that the requirements associated with those skills be discussed and training be held if needed. The skills include:

1. Identification of the two-dozen bad actors – This is done by reviewing all the maintenance records for an asset and using them to identify components with MTBF less than several years. If the normal interval of predictive maintenance needed to deliver 90% reliability is less than quarterly, the component is a candidate for inclusion in MBR.
2. Creating simplified inspections for each of the bad actors – Simplified inspections or health checks are activities that can be used to determine the health of a component simply by close proximity in less than one-minute.
3. Performing the simplified inspections of the bad actors – Executing the simplified inspections requires clearly understanding the characteristics indicative of near-term failure and having the ability to quickly recognize those characteristics.

4. Setting up the “inspection rounds” for the bad actors so the inspections can be done efficiently – Beginning with the objective of performing a complete round of simplified inspections of the two-dozen bad actors in twenty to thirty minutes, the inspections around the asset are laid out in a sequential manner that allows the inspector to move from one inspection to the next within a minute or so.
5. Integrating the MBR activities with the current organizational rhythm – This activity begins with a clear understanding of the work schedule for all of the human resources that will be involved in the MBR process including Operating Supervisor, Maintenance Supervisor, Reliability Engineer, operators, maintenance personnel, assets and the organizations having requirements for those resources. The next step is to identify key elements of the MBR process then, accounting for other requirements, identify the optimum pattern and time at which the MBR process steps will be completed (and can be done well into the future).
6. Gathering inputs on changes occurring with components – Begin by gathering the kinds of symptoms and signs of deterioration and how they will be detected. For instance, a deteriorated bearing is indicated by a hot bearing housing and a bearing housing is considered hot when you cannot hold your hand on it for an indefinite period of time. This symptom would lead to a standard practice of checking bearing condition by attempting to lay an opened hand on the bearing housing.
7. Conducting the “Survival Team” meetings – The bi-weekly “survival team” meeting will be expected to require no more than one-half hour. To fit in that time frame, all the participants must be prepared in advance with specific information and knowledge of the current condition of the asset and its related maintenance programs. For Survival Team meetings held on Wednesday, the objective is to set the stage for preparation of the plan and schedule for tasks on Thursday and ultimately performing all the tasks needed to ensure survival on the following Friday. Then, the asset must survive without breakdown until the Friday, two-weeks later.
8. Planning and Scheduling of the tasks needed to ensure survival – On the Thursday following the Survival Team meeting, all the tasks to be completed on Friday will be planned and scheduled to be completed in the absolute minimum time. To do this, if there are multiple tasks that can be completed simultaneously, a crew of several technicians or mechanics will be needed and the work will need to be scheduled using Critical Path Planning methodology.
9. Execution of regular inspection, PM and CM needed to ensure survival – It is expected that all steps of inspections, Predictive Maintenance, Preventive Maintenance and Corrective Maintenance will be completed using the best practices and techniques. Detailed plans will be assembled and the persons performing the work will follow the steps described in those plans. Initial applications of work plans will be closely observed to ensure that craft personnel have sufficient skills to follow instructions and deliver the desired results.

10. Renewal, Refreshment and Continuous Improvement of the MBR process – Over time, the overall MBR process is likely to need attention resulting in renewal or new and improved ways of doing things, refreshment or steps to once again re-establish the excitement and enthusiasm and Continuous Improvement or finding ways to do things in a more effective or efficient manner. It will be important to recognize when these activities are needed then take the steps needed to infuse the process and ensure continued success.

Conclusion

Mission Based Reliability should be viewed as a useful tool for improving near-term reliability and then long-term reliability by repeating the processes contained in MBR over and over.

MBR depends on identifying the relatively small number of components that cause an asset to be viewed as being less reliable than it would have been, had that small number of unreliable components been engineered in a manner consistent with the vast majority of asset components.

Once the small number of unreliable components are identified, simplified inspections or “health-checks” must be created and applied on an interval that mathematically causes each component to function at a much higher apparent reliability. While not actually improving the reliability of this small number of components, proactive replacements are made at times and on intervals that allow the component seem to be far more reliable resulting in far fewer breakdowns of the asset.

A small team of critical individuals are named and, on a regular schedule, meet to identify all the tasks needed to ensure the survival of the asset until the end of a specific mission interval. This team repeats the process over and over as long as there are components having so short a MTBF that short interval inspection is required to deliver the desired level of reliability.

When simplified inspections identify the need for a replacement, the P-F interval of the device is considered when determining the time by which the replacement must be complete.

While assets like planes and trains have no alternative than to remain using a MBR-like process forever because their entire life is spent completing missions, it may become difficult to sustain the intensity required by the MBR process indefinitely in those instances where it is possible to address the inherent weaknesses in some other manner. Even those individuals without engineering experience are frequently able to identify instances when substandard components must be upgraded to more robust components.

In conclusion, MBR is a useful tool for making near-term improvements to the apparent reliability of an asset and to provide a company and its personnel with the opportunity to create a situation where long-term inherent reliability is possible.