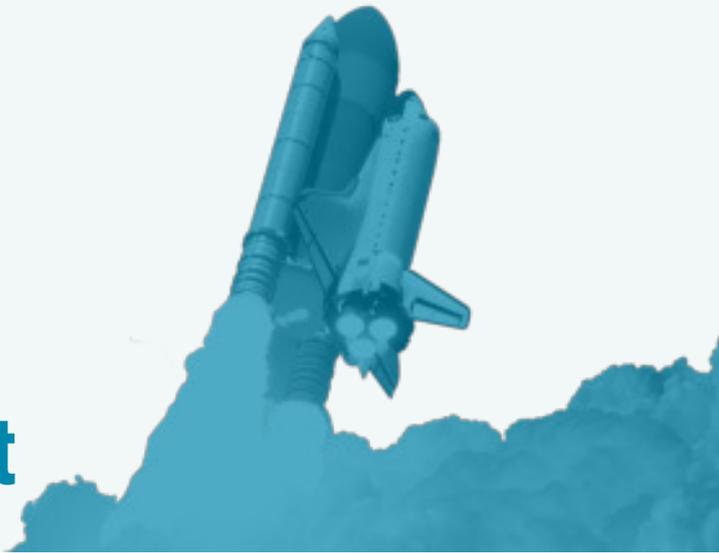


Simple guide to MTBF – What it is and when to use it



Mean Time Between Failure (MTBF) is one of the most widely recognised and yet least understood indicators in the maintenance and reliability world. Manufacturers quote it as a rating of their products and industry uses it as a measure of success. But there is so much misunderstanding associated with MTBF that there is even an online movement to abandon MTBF. In this article, I will explain in simple terms what MTBF is, what it's not, when to use and when not.

What is MTBF?

It is said that the great Greek philosopher Socrates argued that “the beginning of wisdom is the definition of terms.”

Socrates would have been unimpressed with our use of MTBF or would have challenged our collective wisdom when it comes to MTBF.

Sure, there are clear definitions for MTBF. But, unfortunately, there is a lack of common understanding of what MTBF really means.

So, let's start with the definition:

MTBF stands for Mean Time Between Failures and it represents the average time between two failures for a repairable system.

For example, three identical pieces of equipment are put into service and run until they fail. The first system fails after 200 hours, the second after 250 hours and the third after 400 hours. The MTBF of the systems is the average of the three failure times, which is 283.33 hours.

Let's look at some of the definitions of critical terms related to MTBF.

MTBF is related to failure rate. It assumes a constant random failure rate during the useful life of a piece of equipment.

But what do these terms really mean?

We need a clear set of definitions so that we understand what an MTBF number is telling us and what the limitations of that number are. There is even a movement to abandon MTBF because of the misunderstanding and misuse of the term.

We can learn more about MTBF by exploring its origin and the reasons why it came into use. It also helps to compare MTBF with other indicators to avoid confusion about terms. This article covers all these aspects along with some clear guidance about where to use and not to use MTBF.

Failure rate

The failure rate is the number of failures in a component or piece of equipment over a specified period. It is important to note that the measurement excludes maintenance-related outages. These outages are not deemed to be failures and therefore, do not form

part of this calculation. A failure rate does not correlate with online time or availability for operation – it only reflects the rate of failure.

Failure Rate = No. Of Failures / Time

In industrial applications, the failure rate represents past performance based on historical data. But in engineering design, the failure rate can also be predicted. It is common to use a bathtub curve to illustrate failures over the entire life of a product.

There is a high rate of infancy failures at the beginning of its life and a high rate of wear out failures at the end of its life. But in between, during the product's useful life, its rate of

failure is expected to be reasonably constant. Manufacturers seek to reduce infancy failures by testing products and removing early failures before they get to the customer.

The disadvantage of failure rate as an indicator is that it yields a tiny result, which is difficult to interpret. The failure rate of a pump could be 0.4 or even orders of magnitude lower than that.

Reliability

Before World War II, the term reliability described how repeatable a test was. The more repeatable the results, the more reliable the test, whether it be in the field of mechanics, psychology or any other scientific endeavour. However, the challenges of World War II caused new developments in the definitions and engineering associated with reliability.

Electronics equipment during the war was highly problematic. Up to half of the electronic equipment on a naval vessel could be out of service at any time – leading to a renewed focus on understanding and improving equipment reliability. Working groups developed strategies like setting quality and reliability standards for electronic equipment suppliers.

The Advisory Group on the Reliability of Electronic Equipment (AGREE) came up with the classic definition of reliability

"The probability of a product performing without failure a specified function under given conditions for a specified period of time."¹

Around this same time, studies showed that up to 60% of failures in army missile systems were related to component reliability. Military and commercial aviation continued to drive improvements in reliability engineering throughout the twentieth century.

The most commonly used reliability prediction formula is the exponential distribution, which assumes a constant failure rate (i.e. The flat part of the bathtub curve).

Reliability = $e^{-\text{failure rate} \times \text{time}}$

Engineers report reliability as a percentage. It indicates the probability of failure for a piece of equipment in the time given. Reliability does not predict when the equipment could fail during that time, but only the chance of that failure occurring at any point during the time given.²

MTBF

We calculate MTBF by dividing the total running time by the number of failures during a defined period. As such, it is the inverse of the failure rate.

MTBF = running time / no. of failures

During normal operating conditions, the chance of failure is random. It could happen at any time on the flat part of the bathtub curve, just as easily as it could at any other time. Using the exponential distribution for reliability calculation, the MTBF then represents the time by which 63% of the equipment has failed. I.e. Only 37% of components are still in service.

The history of MTBF

The MTBF calculation comes out of the reliability initiatives of the military and commercial aviation industries. It was introduced as a way to set specifications and standards for suppliers to improve the quality of components for use in mission-critical equipment like missiles, rockets and aviation electronics. The military handbook containing MTBF information for electronics Mil-HDBK 217 is discontinued, but other resources like The Telcordia still make use of the military handbook.

Maintenance practitioners first used MTBF as a basis for setting up time-based maintenance strategies. Inspection intervals and routine maintenance tasks were set up based on MTBF. These programs aimed to identify potential failures before they occurred, but time-based systems are not the most effective strategy. Condition monitoring is one example of a strategy that is far more effective for predicting failure than time-based programs based on MTBF.

How to calculate MTBF

As mentioned in the definition, MTBF is calculated by dividing the total time by the number of failures. Let's look at a few examples:

Assuming a situation where there are 1,000 cars that run for one year. If one car fails in that time, the MTBF would be:

$$\text{MTBF} = (1 \text{ yr} \times 1,000 \text{ cars}) / 1 \text{ failure} = 1,000 \text{ years per failure}$$

In an unusual case, consider the MTBF of human life, assuming a population of 500,000. If during the course of a year, 625 people died of random causes, the MTBF would be:

$$\text{MTBF} = (1 \text{ yr} \times 500,000 \text{ people}) / 625 \text{ deaths} = 800 \text{ years per death}$$

This example highlights where MTBF could be misleading as no human being expects to live for 800 years.

In a population of 500 ANSI pumps in water service across multiple sites, 600 fail in a period of three years. The MTBF would be:

$$\text{MTBF} = (3 \text{ yrs} \times 500) / 600 \text{ failures} = 2.5 \text{ years per failure}$$

On their own, these numbers provide some information about reliability but not enough to fully understand the reliability performance of the equipment.

Life expectancy of equipment

Every equipment has a life expectancy based on its components, its design, operating conditions and maintenance history. But not everyone is talking about life expectancy in the same way when they use the term. The service life, the mission life and the useful life of a piece of equipment all refer to different things. We can unpack those differences in more detail.

Service life

Service life refers to the entire duration of an equipment's use. We measure it from the time of commissioning to its final failure or decommissioning.

Engineers also predict service life based on the design specifications. A service life prediction would typically be used in calculations to justify the capital expense of a new asset. Actual service life can be compared with the design service life of a piece of equipment to determine whether it met the expectations of engineers when it was first purchased.

One unique example is that of a missile. By nature, we expect a very high MTBF for a missile indicating the very low probability of failure. But the service life of a missile is very short. It can be as little as a few minutes from the time a missile is fired to the time it explodes.

Mission life

Mission life is the duration used for reliability calculations and analysis. For example, we base the failure rate calculation on the number of failures in a specific time. This time is known as the mission life.

Engineers use reliability indicators to predict failures and make decisions about the future mission life of their equipment. This includes making decisions about spares holding or maintenance strategies for a mission life of the next five years.

Useful life

Useful life refers to the flat part of the bathtub failure curve. It leaves out the time associated with infancy failures at the beginning as well as the time associated with wear out failures at the end of a product's life. Useful life is, therefore, the operational life of any piece of equipment.

In design terms, it reflects the maximum life expectancy of any equipment during normal operations. The useful life does not take into account operating conditions or maintenance history – it assumes a constant and random failure rate.

MTBF versus MTTF

Mean Time To Failure (MTTF) is closely related to MTBF. The difference between the two is that MTTF applies to non-repairable systems, while MTBF applies to repairable systems. In other words, the MTTF calculation is as follows:

MTTF = service time / no. of failures

Engineers determine MTTF by observing a large number of identical components and their combined service time. In this way, it gives some indication of the probability of failure. It is an important indicator for complex systems where some parts cannot be replaced but could impact on the MTBF of the system as a whole.

A fan belt in a motor is a typical example. Fan belts should have an MTTF that is higher than the MTBF of the equipment into which it fits. Otherwise, the whole equipment may fail when the fan belt fails. This correlation provides a key for improving an engineering design. The way to improve MTBF of a complex system may be to purchase better quality parts that have a higher MTTF performance. Nevertheless, one must always bear in mind that MTTF and MTBF are probability related and do not guarantee the life of a piece of equipment up to that duration.

MTBF versus MTTR

Mean Time To Repair (MTTR) describes the average time to execute a repair on the equipment over a given period. It is calculated by adding together the total time for repairs and then dividing by the number of failures during that period.

MTTR = total repair time for all repairs / no. of failures

This acronym could also describe the Mean Time To Recovery, which is slightly different. When using recovery as the basis, the time added must include the notification time of maintenance tasks. In other words, besides the repair time, there is additional time to diagnose the fault and plan the repair. Using recovery as the basis for the calculation gives a higher result than using repair time alone.

MTTR does not give enough information on its own to improve maintenance performance. Reasons for the duration must be investigated to determine whether the time to repair can be reduced. Strategies to reduce repair times

may include spares holding strategies or developing in-house skills instead of relying on outside contractors.

Lengthy repairs have the potential to cause a loss in production. Where this is the case, the losses are usually much more significant than the cost of the repair itself. Loss of production adds a significant economic incentive to minimise the MTTR of mission-critical equipment.

MTTR is different to MTBF. Having both results available gives more information to engineers than either one gives on its own. Equipment that fails regularly but is quick to repair needs a different reliability solution to equipment that hardly ever fails but takes a long time to repair.

What is reliability prediction?

Reliability prediction is an attempt to estimate the failure rate of a complex product made up of several components. It comes from the field of electronics, and this is where it is most often applied.

Electronics manufacturers use empirical handbooks for reliability prediction using MTBF. These books offer predicted MTBF for different electronic components based on field failure rates with some simplifying assumptions. But the handbooks are usually conservative in their estimates and ignore differences in the application design, which could influence failure rate significantly.

Manufacturers use the component MTBF data to calculate an estimated MTBF of their product made up of multiple components – this is known as reliability prediction.

But the limitations of using the handbooks and their assumptions must be taken into account when using predicted reliability information. Predicted reliability is most useful for comparative purposes. For example, a manufacturer could compare the predicted MTBF of different components to help them choose the most appropriate component for their product.

There are two main methods of reliability prediction, with one variation included:

- The parts count method uses the failure rate of the various components as well as the count of components to calculate a failure rate for the product itself. It is a theoretical exercise and can only be verified once the product is in service, and an actual failure history is established.
- The parts stress method uses actual field information from large numbers of the component operating within its rated conditions. Engineers use this historical data as a base for predicting the failure rate of products sold in the present. Of course, field information is not available when a new component comes onto the market. Therefore, some manufacturers use a modified version of the parts stress method known as the accelerated life testing method.
- The accelerated life testing method seeks to establish failure statistics for a product by placing it under high stress, for example, operating a component at a higher temperature higher than its rating. These extreme operating conditions cause premature component failure. Engineers use this failure information to back-calculate predicted reliability under normal operating conditions.

Different electronic handbooks use different assumptions and choosing one over the other could lead to considerable differences in MTBF prediction. Comparing MTBF calculations using one set of assumptions with an alternative calculation based on different assumptions is meaningless. On the other hand, using the same base assumptions to compare components or designs is more helpful.

What MTBF is not

There is some opposition to the use of MTBF as a reliability indicator. Proponents of this view have gone to the extent of creating a movement called “nomtbf”. There is a website of that name and several resources that argue that MTBF is not useful as a reliability indicator or even misleading. Let’s consider some of the objections.

1. People commonly mistake MTBF as an expected life of a piece of equipment before failure. The first part of the indicator – “Mean Time” give the impression that on average, each equipment should last at least this long. But MTBF is based on a probability distribution where the expected failure rate is constant. The resultant exponential distribution gives a result of almost 63% failure by the MTBF value. In other words, only 37 % of equipment remain operational by the time they reach their MTBF.
2. In cases of extreme misunderstanding, some people mistake MTBF as the minimum expected time between failures. This mistaken view leads to significant disappointment because 63% of equipment have already failed by then.
3. MTBF offers no information about the cause of failures. Therefore, it does not yield any insights about what could prevent the failure from reoccurring. Only a root cause analysis can deliver this additional and highly valuable information for improving reliability performance. Failures are not random in practice. They are caused by operating conditions that differ from design conditions, the quality of maintenance, the quality of spares used in repairs and human error – to name a few. Eliminating causes of failure is a significant contributor to improving reliability performance, but MTBF does not contribute to that vital process.
4. The same MTBF result can mean very different things from an equipment reliability perspective. For example:
5. If you have 1,000 cars each driving one mile, and one of those cars fails – you get an MTBF of 1,000 by dividing the total miles by the total failures. On the other hand, if you get a single car driving 1,000 miles during which it fails once, you also get an MTBF of 1,000. These are quite different scenarios, and they reflect different reliability performance, but yield the same MTBF.
6. MTBF assumes a random and constant failure rate – the flat portion of the bathtub curve. The assumption is simplistic and does not reflect real-world conditions. Many pieces of equipment have an increasing probability of failure, the longer they operate. A different probability distribution would give a better correlation with real-world conditions and would, therefore, provide more meaningful information from a reliability perspective.

Misunderstanding MTBF can lead to poor business decisions that are costly to organisations. Using MTBF without additional information about the causes of failures and how to predict failures fails to take advantage of the multiple tools for maintenance and reliability available to engineers. Rather than build a maintenance strategy on a theoretical constant rate of failure, maintenance practitioners can build their strategy around current condition monitoring results and predictions of failure.

When not to use MTBF

MTBF should not be used when the bathtub curve does not represent the actual failure rate. If the component has a wearing part, which increases the chance of failure over time, then MTBF will not accurately describe the probability of failure. In this case, MTBF over-predicts failures early in the equipment's life and under-predicts failures the later part of its life.

The best approach for deciding whether to use MTBF is to first establish the reasons behind the need for this information. For example, if the need is to set spares holding requirements, then there may be a better approach or more information required to make that decision. If the need is to estimate the expected mission or service life of a piece of equipment, then MTBF is not the right tool for that task.

When to use MTBF

In my opinion, it is not necessary to throw out MTBF completely as a maintenance and reliability indicator. We need to understand its limitations and its benefits and use it as one of many tools that help us improve the reliability of equipment in our area of responsibility. Some ways that we can use MTBF include the following:

MTBF is a great way to compare similar equipment operating in similar conditions in terms of performance. A Waterworld article³ highlights this point. The article quotes an average MTBF of 2.5 years for an ANSI pump. Poor performance for this pump is 1.5 to 2 years MTBF, and excellent performance is more than 4 years.

Maintenance and reliability practitioners can use this information to evaluate the performance of their equipment. If their ANSI pump falls into an acceptable range, they may turn their attention to other equipment that could benefit from more direct intervention. But if their pump is performing poorly, it gives them the motivation to investigate the reasons why and come up with corrective measures.

Another good use of MTBF is to monitor progress in reliability initiatives. It is a lagging indicator meaning that the current MTBF result reflects the effectiveness of past actions. Once a reliability program is implemented – like condition monitoring, risk-based inspection or other RCM strategies, it is crucial to measure the impact of that program.

Over time, equipment should become more reliable, and therefore, MTBF should increase. If there is no noticeable change in MTBF, then the reliability program is not achieving its objectives. A positive trend of MTBF over time for equipment on site gives maintenance and reliability practitioners confidence that their programs are achieving the desired results. However, reliability initiatives may take some time to reflect in the lagging indicators like MTBF.

MTBF is also useful for engineering design. Engineers use MTBF in electronic manufacture to compare the effect of using different components in an electronic product. It also helps identify design weaknesses. There may be one component that lowers the MTBF of the product as a whole, and a single change could make a significant impact on design reliability. Electronic manufacturers choose components that meet their overall MTBF objective. Over-specifying components adds to the cost of the product, but under-specifying could lead to premature failures and customer dissatisfaction.

When using MTBF information for design, it is important to understand the parameters of the manufacturer's claims. If MTBF from one manufacturer covers a broader range of operating conditions, it may not be directly comparable with figures quoted from another source.

Conclusion

In this article, we have explored the idea of MTBF – its origins, the misunderstandings people have about its meaning and the ways it is used and abused.

While there is a movement to abandon the use of MTBF completely, it does serve a purpose when its limitations are understood and when used in conjunction with other information.

MTBF is a helpful tool for comparative purposes. It used to evaluate different design options and make choices about components. During the service life of a piece of equipment, it can be used to compare performance against other similar equipment in similar service. This comparison helps maintenance and reliability practitioners to make wise decisions about where to use their time and energy. Lastly, it can be used as a lagging indicator to evaluate the effectiveness of reliability programs like condition monitoring and risk-based inspection.

References

1. History of Reliability Engineering, James McLinn, American Society for Quality – Reliability and Risk Division, <https://www.asqrd.org/home/history-of-reliability/>
2. Reliability Engineering Principles for the Plant Engineer, Drew Troyer, Reliable Plant, <https://www.reliableplant.com/Read/18693/reliability-engineering-plant>

About the author



Erik Hupjé is the founder of the Road to Reliability™ and has over two decades of experience in the areas of maintenance, reliability, and asset management. Erik has a passion for continuous improvement and keeping things simple. Through the Road to Reliability™, he helps Maintenance & Reliability professionals around the globe improve their plant's reliability and their organisation's bottom line.

Erik worked and lived in the Netherlands, the United Kingdom, the Philippines, and the Sultanate of Oman for multinationals like Shell and ConocoPhillips. He is now based in Brisbane, Australia where he lives with his wife Olga and their three children. Personal interests include traveling, cooking, history, and beach fishing.