

MTBF, MTTR and Weibull: reading reliability metrics correctly

Compute MTBF, MTTR and availability from a hypothetical record, then see why MTBF is not a lifetime and how the Weibull shape parameter β bears on scheduled replacement.

In short

- MTBF is operating time divided by the number of failures, and MTTR is total repair time divided by the number of repairs.
- MTBF is not a lifetime: under a constant failure rate, about 63.2% of units have failed by the time one MTBF has elapsed.
- A Weibull shape parameter above one indicates wear-out, so scheduled replacement can help only when that shape is supported by the data.

01 / Three metrics, three questions

Reliability metrics answer different questions. Mean time between failures (MTBF) describes how often a repairable system fails. Mean time to repair (MTTR) describes how long a failure takes to fix on average. Availability describes the share of time the system is ready for use.

No single number is enough. Frequent short failures pull MTBF down, while one very long repair pulls the MTTR average up. Reading the three together shows which loss is actually large: failure frequency, or repair duration. A plant that repairs quickly but fails often needs a different response from one that rarely fails but waits a long time for parts.

02 / Formulas and a hypothetical record

For a repairable system, MTBF equals operating time divided by the number of failures. MTTR equals total repair time divided by the number of repairs. Availability is $A = \text{MTBF} / (\text{MTBF} + \text{MTTR})$.

Consider a hypothetical pump record: 8,000 operating hours, 4 failures, and 24 repair hours across those 4 repairs.

- $\text{MTBF} = 8,000 / 4 = 2,000$ hours
- $\text{MTTR} = 24 / 4 = 6$ hours
- $A = 2,000 / (2,000 + 6) = 2,000 / 2,006 \approx 0.997$

For this record, availability is about 99.7%. It describes a past period, and it holds in the future only if operating conditions stay similar. When comparing several units, use the same operating-time basis for each, or the averages will not be comparable.

03 / MTBF is not a lifetime

The most common mistake is to read MTBF as the average life of a part. MTBF summarises the average time between failures; it does not tell you when a specific component will fail. Life distributions are often skewed, so the typical life can sit well away from the average.

Under a constant failure rate, that is, an exponential distribution, the chance of surviving one MTBF is $e^{-1} \approx 0.368$. Equivalently, about 63.2% of units have failed by the end of one MTBF.

MTBF also has practical limits. It assumes operating conditions stayed the same, and with few failures its confidence interval is wide: 4 failures give only a rough estimate. Treat the number as a summary of a particular period, not as a forecast.

04 / Weibull shape and scale

The Weibull distribution describes age-dependent failures with two parameters. The shape parameter β shows how the failure rate changes with age. When β is below 1, early-life failures dominate. When β is about 1, failures arrive at random. When β is above 1, wear-out dominates.

The scale parameter η is the characteristic life: the age by which 63.2% of units have failed. This follows from $F(\eta) = 1 - e^{-1} \approx 0.632$. The bathtub curve combines the three phases into one picture of a population's failure rate over time. A Weibull fit needs more than failure times, because units still running at the end of observation also carry information; they enter the model as suspensions.

05 / A hypothetical Weibull calculation

Take a hypothetical component with $\beta = 2$ and $\eta = 1,000$ hours. Because β is above 1, the failure rate rises with age, which indicates wear-out.

- Failed by 500 hours: $F(500) = 1 - e^{-(500/1,000)^2} = 1 - e^{-0.25} = 1 - 0.7788 = 0.221$
- Failed by 1,000 hours: $F(1,000) = 1 - e^{-1} = 0.632$

About 22% of these hypothetical components have failed by 500 hours, and about 63% by 1,000 hours. The numbers describe the example, not any real part. For a wear-out component, a planned change can be weighed against the cost of an unplanned failure.

06 / Where to start on your records

Start by structuring the failure records per asset class. For each event, keep the failure time, operating hours, repair duration and the unit's status: still running, replaced before failure, or failed. The status field matters, because units that were replaced or are still running must not be counted as failures.

Next, calculate MTBF and MTTR separately, and plot the distribution of repair times rather than relying on a single average. Fit a Weibull model to the failure data and report β together with its uncertainty. Consider scheduled replacement only when β is clearly above 1 and the planned replacement cost is lower than the expected cost of an unplanned failure. That comparison should involve maintenance and production planning together.

Try it

In horizon, open the built-in example data, find one asset's threshold-crossing time, and write a sentence on why that time answers a different question from an MTBF value.

References and online guide

[Open the online guide](#)

[Mean time between failures \(Wikipedia\)](#)

[Mean time to repair \(Wikipedia\)](#)

[Weibull distribution \(Wikipedia\)](#)

[NIST/SEMATECH e-Handbook: 8.1.6.2. Weibull](#)

Proposed pilot steps do not claim field accuracy. Synthetic examples are not measurements from a production site.