

Remaining useful life (RUL): report an interval, not one number

Learn how RUL is defined, what each of three common approaches needs as data, and how to report P10, P50 and P90 values with a worked trend calculation.

In short

- RUL is the time until a health indicator crosses its failure threshold, and it should be reported as an interval rather than one number.
- When the degradation rate is uncertain, a trend-based forecast gives a wider band, and that band grows as the forecast horizon increases.
- Late predictions usually cost more than early ones, so the failure threshold and the error asymmetry should be written down explicitly.

01 / What remaining useful life means

Remaining useful life (RUL) is the time an asset can still perform its job acceptably, counted from now. In practice it is usually defined through a health indicator (HI): a vibration amplitude, a temperature difference, or a score derived from several signals. The HI trends in the direction of degradation over time. When it reaches a failure threshold, the asset is treated as no longer able to do its job.

The definition raises two requirements. First, the threshold must be written down. Is it a permitted upper limit, or the point where the function is lost? Second, RUL is a forecast, and the future is uncertain. For that reason a prediction should be expressed as a probability distribution, not as one number of days.

02 / Three approaches and their data needs

There are three common families. Life-data methods fit an age-based distribution, such as the Weibull distribution, to the failure times of similar units. They do not know the current condition of one machine; they describe the population.

Degradation-trend methods extrapolate the health indicator forward. They work when the indicator shows a clear and consistent trend. Similarity and machine-learning methods are trained on comparable assets that were run to failure, and they estimate the remaining life of a new unit by matching its pattern to those histories. Each family needs different data: age records, a stable trend, or a set of run-to-failure histories.

03 / A hypothetical trend calculation

Consider a hypothetical pump. Its health indicator is 0.6 today, the failure threshold is 1.0, and the indicator rises by 0.2 per week. The time to threshold is:

- Remaining life = $(1.0 - 0.6) / 0.2 = 2$ weeks

The rate is not known exactly. Suppose it lies between 0.15 and 0.30 per week. The shortest estimate uses the fastest rate: $0.4 / 0.30 = 1.33$ weeks. The longest uses the slowest rate: $0.4 / 0.15 = 2.67$ weeks. Because a higher rate shortens the remaining life, the upper bound of the rate gives the shortest life.

The calculation assumes a linear rise. If the indicator is accelerating, the real remaining life is shorter than 2 weeks, and the linear estimate overstates it. Check the residuals of the trend fit before relying on that assumption.

04 / Reading P10, P50 and P90

Uncertainty comes from several sources: error in the estimated rate, measurement noise, and changes in future operating conditions. Together they form a distribution of RUL, and reported values are percentiles of that distribution. P10 is the value below which 10% of outcomes fall, so it is the early end. P50 is the median: half of the outcomes are shorter and half are longer. P90 is the late end.

The band widens as the forecast horizon grows, because a small rate error becomes a large error over a long projection. For planning, the early end often matters more than the median. A late prediction can turn a planned replacement into an unplanned stop, while an early one usually means a part was changed sooner than necessary. That cost difference should be written into the decision rule before the threshold is set.

05 / Limits and failure modes

The threshold is the most common source of error. It is an engineering decision, not a statistical output, and a model cannot infer it from data alone. A good trend fit to a wrong threshold still gives a wrong life.

Regime changes matter too. A new load profile, a different raw material or a change of operating mode can alter the degradation rate. A maintenance reset drops the indicator suddenly, so trend data from before the reset do not describe the period after it. Finally, with few failures in the history, the fitted distribution is weak. An average from a handful of run-to-failure records is not a reliable population model.

06 / How to start on your plant

Start with one asset class and one health indicator that operators already understand. Write the failure threshold as a loss of function and have the maintenance lead sign it. Then collect the maintenance history: failure dates, what was replaced, and the operating load and speed.

Check whether the indicator trends before building any model. If it stays flat until failure, the method has no signal to extrapolate. Run the forecast alongside the existing maintenance process for a period. Record each prediction, the interval it gave, and when the event actually happened. Judge the method by its early-warning behaviour and its false alarms, not by the central value alone.

Try it

In horizon, open the built-in example data, note the threshold-crossing time with its P10, P50 and P90 values, and write which of the three you would use for a maintenance decision and why.

References and online guide

[Open the online guide](#)

[Prognostics \(Wikipedia\)](#)

[Predictive maintenance \(Wikipedia\)](#)

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

Reliability engineering (Wikipedia)

NASA Prognostics Center of Excellence: data set repository

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