OpsPilot

Weibull Analysis — User Manual

Failure Data & Optimal PM Intervals · AI Engineering Co-Pilot

What this guide covers — what Weibull analysis tells you, how the OpsPilot module collects and interprets failure data, the data you need (including the one input most people forget), and the report you receive.

1. What is Weibull analysis?

Weibull analysis reads a set of failure data and tells you the pattern of failure over an asset's life — and from that, the optimal preventive-maintenance interval. It answers the question maintenance planning keeps guessing at: is replacing this part on a fixed schedule actually doing any good, and if so, at what age?

The whole method turns on three numbers, which OpsPilot explains in plain English:

Parameter	Meaning
Beta (β) — shape	How failures are distributed over life. β < 1 = infant mortality (early failures); β = 1 = random failures (age-independent); β > 1 = wear-out (failures increase with age).
Eta (η) — scale	Characteristic life — the age by which 63.2% of the population has failed.
B10 life	The age at which 10% of the population has failed — a common design and replacement benchmark.

Why β matters most: if β is around 1 (random failures), time-based replacement does nothing — you are throwing away good parts and not preventing failures. Scheduled replacement only helps when β > 1, i.e. there is genuine wear-out. Weibull is what tells you which world you are in.

2. What the OpsPilot module does

Role	Responsibility
🎯 **AI Coach (OpsPilot)**	Guides correct data collection, interprets the Weibull parameters, calculates the optimal PM interval, and explains what the numbers mean for your decision — no statistics background required.
👷 **Data Owner (you)**	Provides the actual failure data — times to failure AND suspension times (units removed from service without failing). Suspensions are critical and the most commonly missed input.

3. The data you need — including the one people forget

- **Times to failure.** The age (hours, cycles, or starts) of each unit when it failed.
- **Suspension (censored) times.** The age of units that were removed, replaced or are still running WITHOUT having failed. Leaving these out biases the result badly — it makes equipment look far less reliable than it is. This is the single most important thing to get right.
- **The specific equipment and failure mode.** Weibull is per failure mode — mixing modes invalidates the analysis.

4. How it works

- Define the equipment and the specific failure mode being analysed.
- Collect the times to failure and — equally — the suspension times.
- OpsPilot fits the distribution and reports β, η and B10.
- It interprets the pattern (infant mortality, random, or wear-out) in plain terms.
- Where wear-out is present, it calculates the optimal PM interval that balances the cost of planned replacement against the cost of failure.

5. What you receive — the output

A Weibull analysis summary report (Word) containing:

- The equipment and failure mode analysed, and the data set used (failures and suspensions).
- The fitted parameters — β, η and B10 — with a plain-English interpretation.
- The failure pattern and what it means for maintenance strategy.
- The optimal PM interval recommendation where applicable, or a clear statement that time-based replacement is not justified.

6. Worked example (illustrative)

Suppose a set of pump bearings returns β ≈ 2.4 and η ≈ 9,000 hours. β well above 1 means genuine wear-out — failures climb with age — so a time-based replacement is worth doing, and OpsPilot calculates the interval that minimises total cost (typically a fraction of η, below the B10 life). Now suppose instead the same bearings returned β ≈ 1.0: failures are random, age tells you nothing, and scheduled replacement would waste good bearings without preventing failures — the report would say so plainly and point you toward condition monitoring instead.

7. Getting the best result

- **Never omit suspensions.** Including only failures is the classic Weibull mistake and it skews everything.
- **One failure mode at a time.** Separate seal failures from bearing failures — don't pool them.
- **Read β before anything else.** It decides whether time-based maintenance even makes sense.

- **More data, more confidence.** A handful of points gives a direction; a fuller data set gives a decision.