PM vs Run-to-Failure: The Real Crossover

> Priced without an unplanned-downtime premium, running an asset to failure looks $2,933 a year cheaper than a PM program. Price a realistic 150% premium on an unplanned hour instead, and running to failure costs $6,067 a year more.

A PM program has a visible, scheduled cost: a visit fee, paid on a fixed interval, whether or not the asset was about to fail. Running the same asset to failure instead looks free until it breaks. That asymmetry, one cost you see coming and one you do not, is exactly what makes this decision easy to get backwards if the unplanned-downtime premium is left out of the math.

Two costs that are not the same kind of hour

An hour of planned PM downtime is scheduled. Parts are staged, labor is booked, and production has a workaround ready. An hour of unplanned failure downtime has none of that. The line stops without warning, the fix happens under time pressure, and the cost per hour is typically higher than a planned hour for the exact same asset. Treating a failure hour and a PM hour as financially identical is the single most common way this decision goes wrong.

A worked example: the same asset, with and without the premium

Take an asset on a 30-day PM interval at $800 a visit, running $9,733 a year. Without any PM at all, this asset is expected to fail 4 times a year, each event costing 5 hours of downtime at $300/hour plus $200 in repair cost. Priced with no unplanned-downtime premium (treating a failure hour the same as a PM hour), that is $1,700 per failure, $6,800 a year, which looks cheaper than the $9,733 PM program. A cost comparison stopping there would recommend dropping PM.

Price the same failures with a realistic 150% unplanned-downtime premium instead, and each failure actually costs $3,950, $15,800 a year, well above the PM program's $9,733. The premium alone is what flips the recommendation from "skip PM" to "run PM." Nothing about the asset changed, only whether the comparison accounted for what an unplanned hour really costs.

Why the naive number is so persuasive

The naive, premium-free comparison is not a made-up worst case, it is the number you get from the most obvious way to do this math: multiply hours by a single dollar-per-hour rate and stop. It produces a real number, it just answers the wrong question, because it assumes an unplanned failure costs exactly what a scheduled outage would have cost. Most operations do not track

downtime cost separately by whether it was planned, so this mistake is easy to make without anyone noticing.

Where the crossover actually sits

The premium needed to flip the verdict depends on the gap between the PM program's cost and the naive run-to-failure cost. A PM program that is only slightly more expensive than the naive failure cost flips at a small premium. A PM program that looks expensive by a wide naive margin needs a much larger premium to catch up, which is exactly why it pays to check the real premium on a specific asset rather than assume PM always wins or run-to-failure always wins.

> **Key takeaway.** Never compare a PM program against run-to-failure without pricing the unplanned-downtime premium explicitly. Leaving it out does not remove the cost, it just hides it in a comparison that looks complete.

Assumptions

- Unplanned downtime typically costs more per hour than planned downtime, because there is no time to stage parts, labor, or a production workaround. A 150% premium is used in the worked example, editable per asset.
- Failure rate under each policy is treated as a single point estimate, not a hazard-rate curve.

Sources

- NIST Handbook 135 (equivalent annual cost framing for capital and maintenance screening)

Frequently asked questions

How do I estimate my own unplanned-downtime premium?

Compare the actual cost of a recent unplanned stoppage (lost production, expedited parts, overtime labor to catch up) against what an equivalent planned outage of the same duration would have cost. The ratio between them is your premium. If you have no history to draw on, 100 to 150% is a reasonable starting assumption for most production equipment.

Does this mean PM always wins once the premium is priced in?

No. A low-criticality asset with a cheap replacement, infrequent failures, and low downtime cost can still favor running to failure even with a realistic premium. The point is not that PM always wins, it is that the comparison is not honest until the premium is included either way.

What is the difference between this decision and picking a PM interval?

This decision asks whether to run a PM program on this asset at all. Picking the interval, how often the PM visit happens, is a separate question that only matters once the answer here is yes, covered by the Preventive Maintenance Interval Planner.

Operations Decisions is a set of transparent decision engines for physical operations, not a guarantee of any outcome. This guide is educational.