

101.



Predictive Maintenance 101: From Data to Insights

An introduction to predictive maintenance for manufacturers who are ready to move beyond scheduled maintenance but are not sure where to start. Covers data requirements, model types, and implementation sequencing.

PART 1

Why predict

PART 2

What data + which
model

PART 3

How to start

Every machine is already telling you when it will fail.

The question is whether anyone is listening early enough to act. Predictive maintenance is the practice of listening: watching how each asset behaves, spotting the change that comes before a failure, and scheduling the repair while there is still time to plan it.

Most plants already run a mix of strategies. Some equipment runs until it breaks. Some is serviced on a fixed calendar. A few critical assets get inspections and spot checks. Each approach has a place, but scheduled maintenance carries a hidden cost: it treats every machine as if it wears at the same rate. Some parts are replaced with life left in them. Others fail between intervals.

This guide is written for maintenance, reliability and operations leaders who want to move beyond the calendar but are unsure where to begin. It avoids hype about algorithms and focuses on three practical questions: which assets to start with, which data you need, and in what order to build the capability.

How to read this guide

PART 1	Why prediction pays and the two curves that govern it.
PART 2	The data and the model types, with a decision tree for choosing between them.
PART 3	A 90-day sequence, the questions we hear most, and a worksheet for shortlisting your first assets.

KEY TERMS

PREDICTIVE MAINTENANCE

Maintenance triggered by data that forecasts a failure, rather than by a calendar or a breakdown.

P-F INTERVAL

The time between the first detectable sign of failure (P) and the failure itself (F).

BASELINE

A record of how an asset behaves when it is healthy, across normal operating conditions.

ANOMALY

Behavior that departs from the baseline by more than normal variation.

RUL

Remaining useful life: an estimate of how long an asset can run before it fails.

CMMS

Computerized maintenance management system, where work orders and failure records live.

PART 1

01


Why predict when you can schedule?

- What does a scheduled strategy actually cost?
- How much warning does a failure give?
- Which assets justify the effort?

From breakdowns to forecasts

Most plants run several maintenance strategies at once. Each rung of the ladder changes what triggers the work.

FIG. 1 The maintenance maturity ladder



A path, not a switch

Few plants jump from reactive to predictive. Each rung builds the data and habits the next one needs: preventive schedules create a maintenance history, condition monitoring creates sensor baselines, and both feed predictive models.


The right rung for each asset

Run-to-failure is still right for cheap, redundant equipment. Predictive maintenance earns its cost on critical assets where a failure stops a line, damages product or takes a long time to repair.

Two curves every maintenance leader should know


One shows how much warning a failure gives. The other shows why doing more maintenance is not always better.

FIG. 2 The P-F curve



Most mechanical failures develop over time. A bearing starts to wear, vibration rises, temperature follows, and eventually the asset fails. The earlier a signal detects the change at **P**, the longer the window before failure at **F**.

FIG. 3 The cost of maintenance



Too little maintenance and failures dominate. Too much and labor and parts dominate. Scheduled programs guess where the low point sits for an average machine. Predictive maintenance finds it for each machine.




Intervene late enough to use the component's life, and early enough to avoid an unplanned stop.

Not every asset deserves a model

The best first targets fail expensively and slowly. They give warning, and the warning is worth acting on.

FIG. 4 Plotting a typical machining plant



Illustrative positions. Plot your own assets using the worksheet on page 16.

ASK OF EVERY ASSET

How much does an unplanned failure cost in downtime, scrap and repair?

Does the failure develop over time with symptoms you can measure?

Is there enough history, or enough similar machines, to learn what normal looks like?

START NARROW

Pick three to five assets that share a failure mode, such as spindle bearings across several machining centers. One failure mode means one set of signals, one model and one workflow to prove.

PART 2

02


What data, and which model?

- Which signals reveal which failures?
- How much history do you really need?
- When is a simple rule better than machine learning?

Three kinds of data, one asset

Condition signals show how a component behaves. Operating context explains why. Maintenance history shows what failure looked like.

FIG. 5 Anatomy of a monitored asset



Context matters as much as the sensor. A spindle running a heavy cut vibrates more than one running a finishing pass. Without the part and program attached, a model cannot tell normal load from early wear.

Start recording failures consistently today. Incomplete CMMS history is normal. Consistent failure codes from now on are what make advanced models possible later.

Which signal reveals which failure

Match the signal to the failure mode you want to catch, then capture it at the rate that failure demands.

FIG. 6 Common signals and what they reveal

SIGNAL	WHAT IT REVEALS	TYPICAL CAPTURE
Vibration	Bearing wear, imbalance, misalignment, looseness	High frequency, from accelerometers
Temperature	Friction, lubrication loss, overload, cooling faults	Low frequency, continuous
Motor current + power	Mechanical load changes, electrical faults, tool wear	Medium to high frequency
Pressure + flow	Leaks, blockages, pump and hydraulic wear	Low to medium frequency
Machine state + counts	Cycles, run hours, stops and their context	Event-based, from the controller
Maintenance records	What failed, when, and what was done	From the CMMS or work orders

Sampling follows the failure

Bearing and gear faults appear in fast vibration signals, while overheating and lubrication loss develop slowly. Capture each signal at the rate its failure mode needs, rather than everything at the highest rate.


Use what the controller already has

Machine states, counts, motor current and temperatures are often available from existing PLCs. Start there, and add dedicated sensors only where the failure mode calls for them.

Choose the simplest model that answers the question

Model choice follows from the data you have, not the other way around.

FIG. 7 Which model type fits your situation?



Most plants start on the left of the tree. Thresholds and anomaly detection need little or no failure history, so they produce value first and create the labeled events that later models need.

The next page compares all five model types, including physics-based models for high-value, well-characterized assets.

Five model types at a glance

Each type answers a different question and needs more data than the one before it.

FIG. 8 Predictive maintenance model types

MODEL TYPE	QUESTION IT ANSWERS	DATA IT NEEDS	USE IT WHEN
Rules + thresholds	Is a value outside its safe range?	Live signals and agreed limits	Failure modes are well understood
Statistical anomaly detection	Is this asset behaving unusually?	A baseline of normal operation	Failure history is limited
Machine learning classification	Which failure is developing?	Labeled examples of past failures	Failures recur and are recorded
Remaining useful life (RUL)	How long until it fails?	Run-to-failure histories across many assets	Many similar assets share a failure mode
Physics-based models	What does engineering theory predict?	Design data and operating conditions	Assets are high value and well characterized

Start with rules and anomalies

Thresholds and anomaly detection need little or no failure history, so they can run within weeks. They also generate the labeled events that more advanced models need later.

Graduate as data accumulates

As failures and repairs are recorded against the signals that preceded them, classification and remaining-life models become practical. Move up only when the simpler model stops answering the question.

PART 3

03


How to start, and what to expect

- What happens in the first 90 days?
- What goes wrong, and how do you avoid it?
- Where does Intelycx CORE fit?

The first 90 days

Alerts come before algorithms. A simple warning someone acts on is worth more than a sophisticated model no one trusts.

FIG. 9 A 13-week pilot plan



WEEKS 1-4 Connect three to five critical assets and confirm the data matches what maintenance sees on the floor.	WEEKS 5-9 Learn normal behavior, then switch on thresholds and anomaly alerts, each with an owner and a response.	WEEKS 10-13 Record every finding against the data, tune alerts with technicians, and extend to similar assets.
--	---	--

Questions we hear most

The obstacles are usually practical, not technical.

Q We have almost no failure history. Can we still start?

A Yes. Anomaly detection needs only a baseline of healthy operation. It flags unusual behavior from the start, and every confirmed finding becomes the failure history you were missing.

Q How do we stop alert fatigue?

A Set thresholds with the technicians who will receive them, start conservatively, and review alert quality every few weeks. Measure interventions triggered, not alerts sent.

Q How do we know it is working?

A Track planned interventions that avoided a stop, unplanned downtime on the pilot assets, and the lead time between alert and failure.

Q Do we need new sensors on every machine?

A Rarely at first. Controllers already expose states, counts, current and temperatures on many assets. Add vibration sensors where the failure mode calls for it, such as rotating equipment.

Q Who should own the program?

A Maintenance owns the response, operations owns the schedule impact, and IT or OT owns the data. Name one accountable lead for the pilot.

Q When should we move to machine learning?

A When failures and repairs are recorded consistently against the signals that preceded them, and simpler rules stop answering the question.

What you need, and where CORE fits

YOU NEED TO...	INTELYCX CORE PROVIDES
✓ Connect old and new assets	Intelycx Edge communicates with more than 400 device types, including legacy controllers. Nano Edge brings older equipment online.
✓ Capture high-frequency signals	High- and low-frequency signals captured with ultra-low latency.
✓ Add operating context	Asset identity, timestamps, states, counts and production context in a unified namespace.
✓ Alert the right person	Real-time alerts, notifications and P911 for urgent interventions.
✓ Predict at the edge	On-device machine learning raises defect and vibration alerts, even without a cloud connection.
✓ Close the loop	Descriptive, diagnostic and predictive analytics, integrated workflows and automated audit tracking.

CASE IN BRIEF · MEDICAL DEVICE OEM

20%

● BEFORE

Weekly manual OEE collection. Critical CNC machining centers failed without warning.

● DEPLOY

CORE replaced the weekly cycle with a continuous, automated dashboard and predictive maintenance alerts.

● SIX MONTHS

20% reduction in machinery downtime.

Your first-asset shortlist

List candidates, then circle the three to five that score high on cost of failure and warning before failure.

FIG. 10 Shortlisting worksheet

ASSET	FAILURE MODE	COST OF FAILURE	WARNING?	SIGNALS AVAILABLE	HISTORY?
<i>Spindle, MC-04</i>	<i>Bearing wear</i>	<i>High: line stops</i>	<i>Yes, weeks</i>	<i>Current, temp; add vibration</i>	<i>Partial</i>

First row is an example.

Bring this sheet to a working session.

We will help you score the shortlist and plan a 90-day pilot on Intelycx CORE.

intelycx.com →



CONNECTING THE PHYSICAL AND DIGITAL WORLD

The World of Intelycx

Discover the intelligence behind a connected enterprise.

Accelerating performance.
Every minute of every day.

intelycx.com