TECHNICAL NOTE

Vibration analysis: what it can and cannot tell you about a machine

Vibration monitoring is heavily promised and rarely explained. This note sets out the technical conditions that make a measurement reliable, the cases where it is not, and the questions to ask a supplier before signing.

Bludocknet Srl · September 2026

Anyone assessing a predictive maintenance system meets more or less the same sales material every time: a sensor, a dashboard, and the promise of knowing in advance which component will fail. The promise is technically possible, but only under specific conditions, and on a production floor those conditions are often absent.

The outcome is predictable. The system goes in, produces alerts nobody can interpret, and within a few months nobody opens the dashboard any more. Not because the technology does not work, but because it was applied to a machine that did not meet the prerequisites.

This note is written for the person who has to make the call. It contains no performance figures and no customer stories: it sets out the physical constraints of the problem, which are the same for every supplier.

What an accelerometer actually measures

An accelerometer mounted on a housing measures the acceleration of that point along one or more axes. It does not measure the condition of a bearing, a gearbox or a belt: it measures the movement of a piece of metal that receives, summed together, every force generated by the machine and by whatever surrounds it.

From that signal you normally derive an overall severity value, typically RMS velocity in the 10-1000 Hz band as defined by the ISO 20816 series (which superseded ISO 10816), and a frequency spectrum obtained by transform. The overall value tells you *how much* it vibrates. The spectrum is where you are supposed to work out *why*.

Two different routes, often confused

Two very different approaches live under the label “predictive maintenance”, and confusing them is the root of most disappointments.

Component diagnosis

This is the classical vibration analysis approach. You know the geometry and kinematics of the machine, you calculate the characteristic fault frequencies, and you look for them in the spectrum. A rolling element bearing has four frequencies known in the literature as outer race,

inner race, rolling element and cage defect frequencies, which depend on the number and diameter of the rolling elements, the pitch diameter and the contact angle. A gear generates a harmonic at the product of tooth count and rotational speed. Unbalance shows at the first order, misalignment typically at the second.

If those lines appear and grow, you can say which component is degrading. That is genuine diagnosis.

Anomaly detection

This is a different approach, and it does not require knowledge of the machine. You record how it vibrates when it is running well, you build a reference, and you flag when behaviour departs from that reference in a statistically significant way.

It does not tell you which component. It tells you that something has changed, when it changed, and on which axis or frequency band.

> **THE POINT THAT MATTERS**
>
> Component diagnosis is more informative, but its prerequisites are strict. Anomaly detection is less informative, but it survives real conditions. A supplier promising the first on a machine that does not meet the prerequisites either does not know them or is hiding them.

Seven constraints that decide whether a measurement is reliable

1. Constant rotational speed

Frequency analysis assumes the machine always runs at the same speed. If speed varies, every line in the spectrum shifts and smears, and characteristic frequencies become unreadable. There is a remedy: acquire a reference signal from the rotation and move to order analysis instead of frequency analysis. But you need that signal, and on an existing machine it is often unavailable without additional work.

Many modern production machines run under inverter control, with continuous ramps and speed changes. This is the first reason why diagnostic analysis transplanted from the turbomachinery world does not work on the shop floor.

2. Repeatable load conditions

A machine running different products vibrates differently, and legitimately so. A change of format, tooling, recipe or material moves the reference. If the system cannot tell a changeover from degradation, it raises an alarm every time production switches article, and within weeks the operator stops looking at them.

The answer is to build a reference for each operating condition, not a single one. Which means being able to read the state of the machine, not just its vibration: this is why vibration data on its own, without process context, delivers far less than it could.

3. Sensor mounting

This is the most underestimated constraint and the easiest to get wrong. How the sensor is fixed determines the usable frequency band, and the difference between one mounting and another is an order of magnitude:

MOUNTING	USABLE BAND	NOTES
Threaded stud on a spot-faced surface	The widest	Quality benchmark, requires machining
Rigid adhesive bonding	Good	Common compromise, depends on the adhesive
Magnetic base	Reduced	Convenient, but cuts high frequencies
Hand-held probe	Very reduced	Coarse comparative readings only

Early-stage bearing faults show up precisely at high frequencies. A sensor on a magnetic base will struggle to see them until they are already obvious, and by then the warning time has shrunk.

4. Measurement point

The sensor belongs as close as possible to the bearing housing you want to watch, along the stiffest path to it. Every joint, cover or interposed sheet attenuates the signal and changes its content. An accelerometer mounted on a guard mostly measures the resonance of the guard.

5. Spectral resolution against battery life

Separating two close lines in the spectrum requires a long acquisition window: frequency resolution is the inverse of acquisition duration. Long acquisitions are needed to separate, for example, the rotational frequency from the sidebands that indicate modulation.

On a battery-powered sensor this conflicts directly with battery life: the more you acquire and transmit, the shorter it lasts. This is an engineering trade-off, not a commercial detail, and it should be settled before installation rather than after.

6. Shop floor background noise

A rigid machine bolted to a shared floor slab receives vibration from neighbouring machines. In a dense department, part of what the sensor records does not belong to the machine it is mounted on. This does not prevent monitoring, but it raises the threshold below which nothing can be distinguished.

7. The data needed for diagnosis often does not exist

Calculating a bearing's fault frequencies requires its geometric data. On a fifteen-year-old machine, possibly rebuilt, with a bearing replaced by a technician who fitted an equivalent from another brand, that data is gone. No amount of processing can recover it from the signal.

This is why component-level diagnosis is realistic on documented critical assets, and far less so on a heterogeneous machine fleet.

When it is not worth it anyway: the redundant machine

There is one condition that makes monitoring pointless even when everything else is in order, and it is not a technical one.

If a machine stops and production redistributes across other machines with no significant loss, advance warning is worth nothing. Knowing ten days ahead that a motor will degrade changes no decision, because the decision — stop it and replace it — can be taken when it happens anyway.

Condition monitoring has economic value where a machine stopping stops something downstream: a synchronous line, a continuous process, a bottleneck. That is the first question to ask, before any technical assessment.

Questions to ask a supplier

None of these is a hostile question. They are the questions a serious supplier expects, and answers gladly.

- **Does the system do component diagnosis or anomaly detection?** If the answer is the former, ask which kinematic data is required and who provides it.
- **How is variable speed handled?** If the machine runs under inverter control, you need to know whether a rotation reference exists or whether only overall indicators are used.
- **What happens at changeover?** A single reference on a machine running different articles will produce false alarms.
- **What mounting is planned, and what frequency band does it guarantee?**
- **How often is the signal acquired, and for how long?** And with what battery life if it is battery powered.
- **How long is the learning phase before alarms become trustworthy?**
- **Who receives the alarm and what do they do?** An alarm landing in a mailbox nobody reads is a cost, not a service.
- **Where does the data sit and who owns it?** Ask at the start, because once the contract is signed it is no longer negotiable.

In short

Vibration analysis is a mature and reliable tool within its boundary. That boundary is narrower than it is usually described, and the fastest way to waste an investment is to apply it where the prerequisites are absent.

On a critical, non-redundant machine with repeatable operating conditions and a well-mounted sensor, you get warning time that has value. On a heterogeneous, variable-speed machine fleet, the realistic goal is not knowing which bearing will fail: it is noticing that a machine has started behaving differently from how it used to, and going to look.

It is a less spectacular goal. It is also the only one that holds up over time.

Where this note comes from

It was written by Bludocknet, developer of Box4AI, an industrial monitoring platform. Our vibration monitoring module works by anomaly detection: it learns how a machine vibrates when it is running well and flags when that behaviour changes. It does not perform component-level diagnosis, and above we have explained why in many real settings it would not be honest to claim otherwise.

If you are assessing a project and want a view on whether your machine meets the prerequisites, write to us: info@bludocknet.com