Oil Analysis Limits Belong to the Compartment, Not the Lab

The same iron reading is Normal in one compartment and Critical in the next, on the same machine. A lab-wide limit table gets one of them wrong every time.

Brandon Holland, Marketing & Operations Digital Solutions Architect

September 7, 2026 · 10 min read

IN BRIEF

Oil analysis alarm limits should belong to the compartment, not to the laboratory, because an engine, a hydraulic system and a final drive carry different oils and fail differently. Trends should be read within an oil fill, since a drain resets wear metals, and the software's severity should be a proposal the analyst can overrule, on the record.

Key takeaways

- A limit set for the lab produces false alarms in one compartment and silence in another, on the same machine.
- The reading that has not breached but is climbing sharply is the one that predicts the failure. A limit table cannot see it.
- Compare against the same point in the previous oil fill, never across a drain, or every oil change reads as a cure.
- A diagnosis names one fault and treats a consequence as corroboration, not as a second fault to send someone chasing.
- The analyst overrules the software, the record says so, and that record is what should drive the next revision of the limits.

An oil analysis report is judged on its last paragraph. The numbers above it are the same numbers any accredited lab would produce; the value the customer pays for is the sentence that says what is wrong and what to do. That sentence depends on three decisions that most software gets wrong: which limits a reading is judged against, which earlier sample it is compared with, and who decides what the result means.

This article is about those three decisions. It argues that alarm limits belong to the compartment rather than the lab, that a trend must restart at every oil change or it lies, that a diagnosis should name one fault rather than list the numbers, and that the analyst, not the screen, gets the last word. It is written for the used oil analysis lab and the maintenance manager who reads its reports, and it covers how the numbers should be read, not how they are measured.

Why do labs analyse machines rather than samples?

Used oil analysis labs do not analyse samples, they analyse machines. A single sample from a haul truck's engine tells you almost nothing, because the amount of iron in used oil depends on the engine, its age, its duty, how long the oil has been in it and what the last sample showed. The same reading is unremarkable in one context and alarming in another. The information is in the run of samples from the same compartment over months and years, read as a trend, ending in a judgement about what is wrong and what to do.

That shape does not fit a system built around the sample. A sample-centred system shows one record at a time, compares each result to one limit, and has no idea that the record beside it came from the same compartment three hundred hours earlier. It holds the data well and judges none of it. The diagnostician does the joining in their head, or in a spreadsheet, and the value the lab sells is produced outside the system that stores its data.

So the unit of record in an oil lab is the compartment: an engine, a transmission, a hydraulic system, a final drive, a transformer's main tank. Each has its own oil, its own sampling interval, its own history and its own limits. Everything below follows from taking that seriously.

Why should limits belong to the compartment?

Because the compartments on one machine are different machines. A diesel engine, a hydraulic system and a final drive on the same haul truck carry different oils at different viscosities, run at different temperatures, wear different metals and fail in different ways. Iron in a final drive comes from gear teeth and is expected in quantities that would condemn an engine. Copper in a hydraulic system may be a cooler core dissolving; in an engine it may be a bearing. A lab-wide table with one caution and one critical limit per element gets one of those compartments wrong every time.

The extreme case is the transformer. A load tap changer switches current under oil and arcs by design, so its oil carries gases that would take a main tank out of service. The same acetylene reading is normal in one compartment and critical in the other, on the same transformer. There is no limit that is right for both, and a lab that applies one produces either false alarms on every tap changer or silence on a failing main tank.

Limits should therefore be set per compartment type at minimum, and refined per compartment as its history accumulates. Published guidance on statistically deriving alarm limits from a population of similar compartments is the right starting point, and the limit set should be a record the lab can edit, so that lowering the caution limit for iron on the hydraulic set takes effect on the next reading without a software change. The limits are the lab's knowledge about its customers' machines, and they should live where the lab can maintain them.

What does "inside the limit and moving" mean?

A limit table only ever tells you about a compartment that has already crossed a line. By the time iron breaches the critical limit, the bearing has been shedding metal for two or three samples, and each of those samples was inside every limit and looked normal on a table. The failure was predictable and the table did not predict it.

So a second rule is needed beside the limits: a reading that is close to its caution limit and has risen sharply against the comparable earlier sample is marginal, even though it has breached nothing. Marginal means inside the limit and moving. It is the rate-of-change rule, and it is the one that catches a failure early enough to schedule a repair rather than recover from a breakdown. Where a low reading is the fault, as with an additive being depleted, the same rule runs the other way: near the floor and falling.

> **NOTE**
>
> A rate-of-change flag is only as good as the sample it compares against, and that is the subject of the next section. A system that compares each sample with whichever sample happens to precede it will raise marginal flags after every oil change and miss the ones that matter. The rule and the comparison have to be designed together.

A flag raised by this rule should say so in words: which reading, how close to which limit, and which earlier sample it was measured against. A flag that cannot be explained is a flag the diagnostician learns to ignore, and once the flags are ignored the screen is decoration.

Why must the trend restart at an oil change?

Wear metals accumulate in the oil. They climb through an oil's life and drop to near nothing when the oil is drained and replaced. So a compartment's raw trend is a sawtooth, not a ramp, and a naive comparison reads every drain as an improvement and every fresh fill as the machine deteriorating again. The machine did not get better because the oil was changed.

The history therefore has to be cut into oil fills, and the comparable earlier sample means one of two things. If there is a previous sample in the same fill, compare against that: same oil, further into its life. If this is the first sample of a fresh fill, compare against the sample at the nearest oil age in the previous fill, so that a reading at two hundred hours on new oil is judged against the reading at two hundred hours on the last oil, not against the gassy, metal-laden sample taken just before the drain. A new fill starts where the oil was renewed or where its recorded age has gone down; either is proof that the fluid is not the fluid that was in it last time.

Charts should follow the same rule. The trend line should break wherever the oil age falls, so a drain is never drawn as a downward slope. Where a compartment is sampled at every oil change, every sample sits at the same oil age, the readings are directly comparable and the line runs straight through, and the chart should say which of the two it is showing. For transformers, where oil is degassed rather than drained, the same applies: a trend read straight through a degassing job says the fault cured itself, the most expensive misreading of a transformer.

Should the report name the fault or list the numbers?

Name the fault. A report that says iron, chromium and the ferrous debris index all breached is a restatement of the table. A report that says a bearing or gear is breaking up, and that the filter should be cut, is what a diagnostician is paid for. Faults show up as combinations of readings moving together, and the diagnosis is recognising the combination.

FAULT	WHAT HAS TO MOVE TOGETHER	WHAT IT POINTS AT
Coolant ingress	Sodium with potassium or boron, plus water or a falling base number, or glycol present at all	Pressure-test the cooling system and the oil cooler. Additive variation does not move three coolant markers at once.
Dirt ingress	Silicon breaching or rising sharply, with aluminium or iron climbing in step	Air filter housing, inlet ducting and clamps. Silicon rising alone is a sealant artefact, not dirt.
Bearing or gear break-up	Ferrous debris over limit and climbing faster than iron, with iron, chromium or nickel raised	Cut a filter or pull the magnetic plug. Debris running ahead of iron means the metal coming off is getting coarser.
Fuel dilution	Fuel dilution over limit, usually with viscosity falling	Injectors, injector seals, fuel pump, and how much the machine idles.
Bearing overlay wear	Lead or tin over limit, or copper with tin rising, and silicon not raised	Oil pressure and bearing clearances. If silicon were up as well, dirt would be doing the cutting.

Where two patterns match one sample, the report should name one primary fault and give the other a sentence. Where the second pattern is a known consequence of the first, it should be worded as corroboration rather than as a second fault. Coolant attacks bearing overlay and eats the additive package; saying “two faults” there would send somebody looking for a second cause that does not exist. And where nothing matches, the report should say plainly what breached and by how much, and claim nothing more.

Who gets the last word?

The analyst. Everything above is arithmetic and pattern-matching, and it should be presented as a proposal: a severity the software suggests, a diagnosis it drafts, drawn on the screen as proposed rather than as decided. The diagnostician reads it against what they know about the customer, the site and the machine’s history, and releases the sample as whatever severity they judge right. Nothing should be released without a person choosing to release it.

When the analyst overrules the proposal, the record should say so. A released sample whose severity differs from what the software proposed should carry a sentence noting the difference, so that nobody reading it later has to compare two fields to find out. Over time that record is the most useful thing the lab has: it shows where the limits are wrong, where the patterns miss, and where the analyst’s judgement is consistently better than the arithmetic, which is exactly what should drive the next revision of the limit sets.

The same respect applies to the smaller fields. If the software fills in a resample interval from the severity and the analyst types their own, the software should stop overwriting it. A screen that overwrites a number somebody has typed is a screen nobody should trust. A reissued diagnosis should keep the earlier one readable, with who changed it and why. The customer’s own view of the report should show the released diagnosis and nothing of the proposal, since the proposal was never the lab’s opinion.

> - Limits belong to the compartment and live on a record the lab can edit
> - Every flag says which earlier sample it was measured against, in words
> - The trend breaks at every oil change or degassing job
> - The diagnosis names one fault and treats a consequence as corroboration
> - The severity is a proposal until a named analyst releases it, and the record shows when they differed

What an oil analysis lab carries in this work

LabLynx has served used oil analysis and analytical laboratories for long enough to have seen a lab-wide limit table produce both kinds of error on the same customer's fleet. The review workspace described below is the shape we build toward, and the boundary is stated plainly.

> **How LabLynx handles this**
>
> LabLynx One should hold each compartment's alarm limits as a record the lab can edit, lay a compartment's whole sampling history out as columns with the oil changes marked, compare each reading against the comparable sample within the same oil fill, raise rate-of-change flags in words, draft a diagnosis that names one fault, and propose a severity that the analyst releases or overrules with the difference recorded. It covers used oil on lubricated machinery and insulating oil on transformers as two disciplines that share a queue and nothing else. It does not measure anything, and it does not replace the diagnostician: the last word is theirs, every time.

Frequently asked questions

How are oil analysis alarm limits set?

Per compartment, not per laboratory. A starting set comes from published statistical guidance applied to a population of similar compartments, and the set is refined as each compartment's own history accumulates. The limits should live on a record the lab can edit, so that tightening a caution limit for one compartment type takes effect on the next reading without a software change.

What is a marginal result in oil analysis?

A reading that is inside its limits but close to the caution limit and rising sharply. It has breached nothing, and it predicts a failure, because a bearing that is going shows up as a rising trend before it crosses a line. Where a low reading is the fault, the same rule applies near the floor and falling.

Why does an oil change reset the trend?

Because wear metals accumulate in the oil and are removed with it, so the raw history is a sawtooth. A comparison that ignores the drain reads every oil change as the machine improving. The comparable earlier sample is the previous one in the same fill, or the sample at the nearest oil age in the previous fill.

Should oil analysis software diagnose the fault?

It should draft a diagnosis and propose a severity, and it should present both as proposals for the analyst to release or overrule. Faults appear as combinations of readings moving together, and recognising the combination is work software can do consistently. Deciding what it means for this customer's machine is the diagnostician's, and the record should show when the two differed.

Can the same reading be normal on one machine and critical on another?

Yes, and on the same machine. A final drive sheds iron from gear teeth as a matter of course. A transformer's tap changer arcs by design and carries gases that would take the main tank out of service. Each compartment has its own oil, duty and failure modes, so each needs its own limits. A lab-wide table is wrong for one of them every time.

Sources and references

1. ASTM D7720 Standard Guide for Statistically Evaluating Measurand Alarm Limits when Using Oil Analysis to Monitor Equipment and Oil for Fitness and Contamination

   ASTM International

2. ISO 4406:2021 Hydraulic fluid power. Fluids. Method for coding the level of contamination by solid particles

   International Organization for Standardization

Brandon Holland

Marketing & Operations Digital Solutions Architect

Brandon Holland is LabLynx's Marketing & Operations Digital Solutions Architect. He builds and runs the digital systems behind the company's public surface, from the website and content architecture to the tooling that keeps product information accurate everywhere it appears. He writes about laboratory informatics from the systems side: how lab records get structured, searched, and kept defensible as a lab scales.