RFT Board user guide

Misblend investigation, root-cause tracking and countermeasure register

Coryton Advanced Fuels

17 November 2026

Contents

1. Getting started

2. The dashboard

3. Raising a misblend record

4. Triage

5. Running the investigation

6. Cause coding and 5 Whys

7. Cost of quality

8. The actions register

9. Sign-off and review

10. The KPI report

11. History and reports

12. Wall display

13. Admin

14. Suggesting an improvement

15. Troubleshooting

Appendix A — what counts against the rate

Appendix B — 1 August 2026, and the imported records

Every screen and every field of the Coryton Right First Time Board, written from the user's side of the screen.

For the daily handful, see the quick guide. For ten minutes on the whole thing, open **Help → Training walkthrough** in the board itself.

1. Getting started

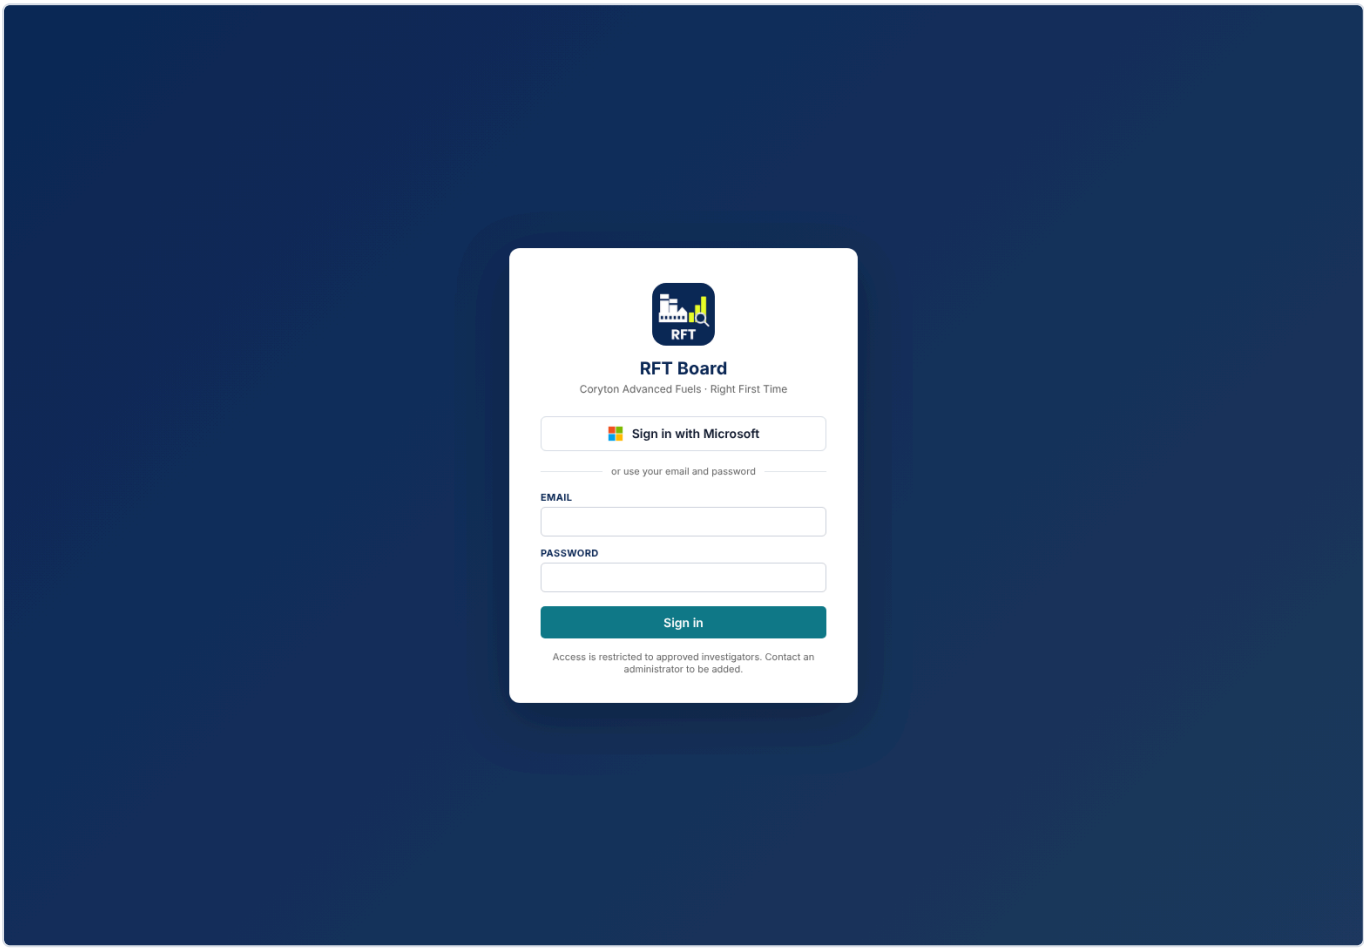
Signing in

The board asks you to sign in before it shows anything. Since 16 November 2026 it does that with your work Microsoft account **by itself**: open the board signed out and it goes to Microsoft without you pressing anything. If you are already signed in to Microsoft on that computer, as you usually are at work, you are back on the board, signed in, a moment later.

The sign-in screen below is what you see when it does not go by itself: after you sign out (signing out keeps you signed out), or when Microsoft says no. It has two ways in, and both open the same account:

- **Sign in with Microsoft**, the button at the top — your work Microsoft account, no separate password.
- Your **email and password**, underneath, as before. To go straight to this form without being sent to Microsoft, add `?signin=password` to the address.

Accounts are made by the **Access app**, when an Access administrator gives you a role on the board; there is no sign-up link on the login screen, and signing in with Microsoft does not make an account. If it says *Your Microsoft account has no account here yet*, ask for access through the Coryton Portal (its **Request access** on the RFT Board tile), or ask an Access administrator. Your name on the board is your name in Microsoft.



If a new version has been deployed since you last used it, the board updates itself at this point before letting you in. There is nothing to lose at the login screen, so it does not ask.

Your role

Three roles, shown as a badge in the header beside the clock.

ROLE	CAN DO
Read only	See every record, report and figure. Change nothing.
Investigator	Everything above, plus raise records, run triage, walk investigation paths, raise actions and prepare reports.
Admin	Everything above, plus the Admin tab — cause codes, triage wording, investigation paths, containment options, product families, hourly rates, blend counts and wall timings. Only an admin can **withdraw** or **delete** a record.

Roles are set in the Access app (<https://coryton-access.netlify.app>), by an Access administrator, not in the board. Access keeps the board in step every 15 minutes: it makes the account, sets the role, and stops somebody set to *No access* from signing in. Since 17 November 2026 nobody changes a role inside the board, an admin included; **Admin → Users & Roles** only lists who has which, because a change made there would be put back by the next sync. An account with no role set in Access reads as **read only**, so it can never change a record before somebody has decided it should.

To ask for a higher level, click **Request more access** beside your role badge (read only and investigator see it; an admin, already at the top, does not). It opens the Coryton Portal in a new tab, so nothing you were typing is lost; the portal signs you in with Microsoft and sends the request, and an administrator decides in the Access app. A read-only user also finds the link on any record they open.

A role change set in Access reaches the board within **15 minutes**, and the person then sees it on their **next page refresh**, not immediately. If somebody says they still cannot edit, check what Access says, give it a quarter of an hour, and ask them to reload before looking for a fault.

A non-admin can open the Admin tab and read all of it — useful for seeing which cause codes and paths exist — but every control is disabled.

Installing it on your device

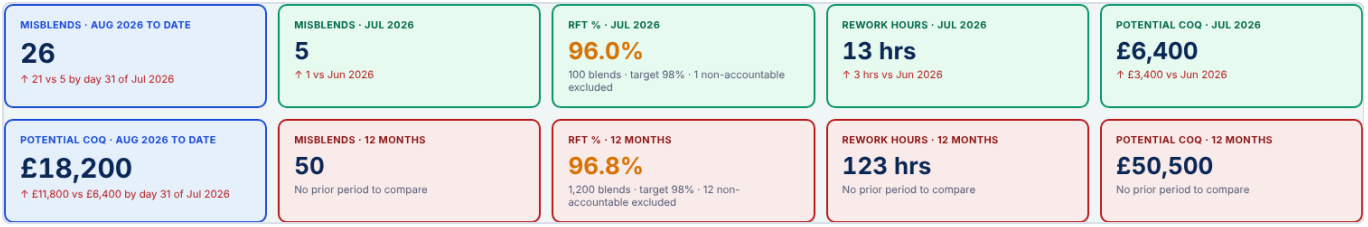
Install the **Coryton Portal** instead: its **Install** button, at the top of the portal, gives it its own icon and its own window, and the board opens from its **Production** or **Technical** tile. The board no longer offers to install itself (Craig's ask, 16 Nov 2026: one app to install, not one per app). If you installed the board before, it keeps working as it did, updates included.

The build badge

The header carries a **Build** badge showing a date. That is the version you are running, and it is the first thing to quote when reporting a problem.

When a new version is deployed mid-session the board shows a banner rather than swapping underneath you. You choose when to apply it, so an update cannot land while you are halfway through editing a record. Dismiss it and it applies on your next full reload.

2. The dashboard



Three periods, three colours

The tiles are three groups. Each is a contiguous block with its own colour and its own window, and the line above the grid names all three so it doubles as the legend.

GROUP	WINDOW
Blue — current month	The month you are in, so far
Green — last complete month	The month just ended
Red — rolling 12 months	The last twelve months that have a blend total

The colour is on the border and the label, never on the number. Red and green already mean bad and good on a figure, so a tinted value would read as a verdict on the number rather than as which group it sits in. A red RFT inside the green group is normal and correct.

Why the current month has two tiles, not four

The blue group shows **Misblends** and **Potential CoQ** and stops there.

RFT % needs a blend count, and this month's does not exist yet — it arrives at month end. A rate tile there read `–` every day of every month, so it was removed. Counts and costs need no denominator and are meaningful from day one, which is why those two stay. The group is drawn boldest because it is the live one.

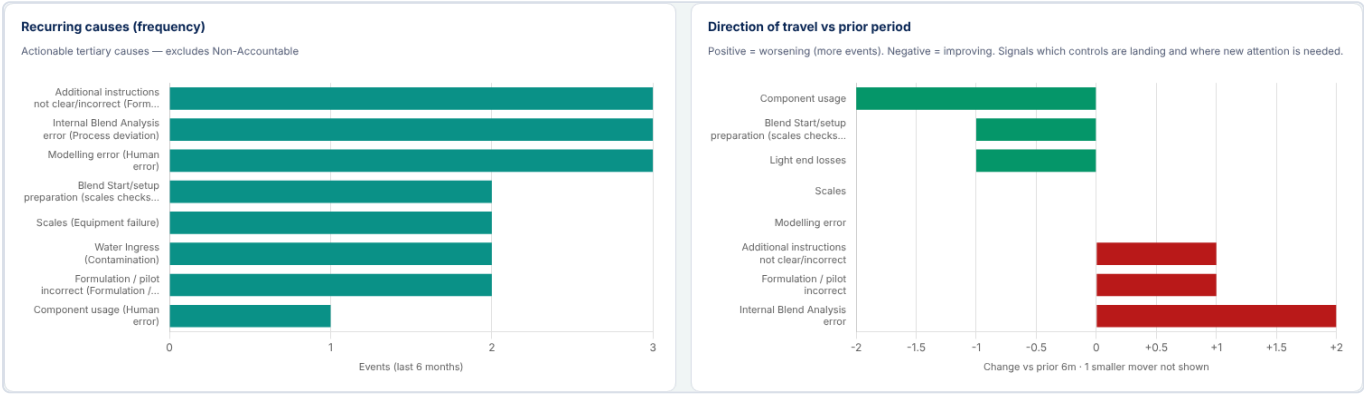
Trend notes

Every tile compares itself against the period before it and says which way it moved. **Up is red and down is green throughout**, because every figure on the dashboard is a count or a cost, so more is always worse.

The current-month tiles compare against **the same day of the previous month**, not its total. A fortnight into the month, comparing against a whole month would always look like an improvement.

A window is never compared against a shorter one. If only six complete months exist before a twelve-month window, the tile carries **no trend note at all** rather than a flattering one.

The charts and the lists



Recurring causes is a Pareto — which causes come up most often. **Direction of travel** compares this period against the one before and keeps the biggest movers. Both keep their own six-month window and say so on the axis: a twelve-month Pareto buries a cause that started in the spring under one that was fixed last autumn.

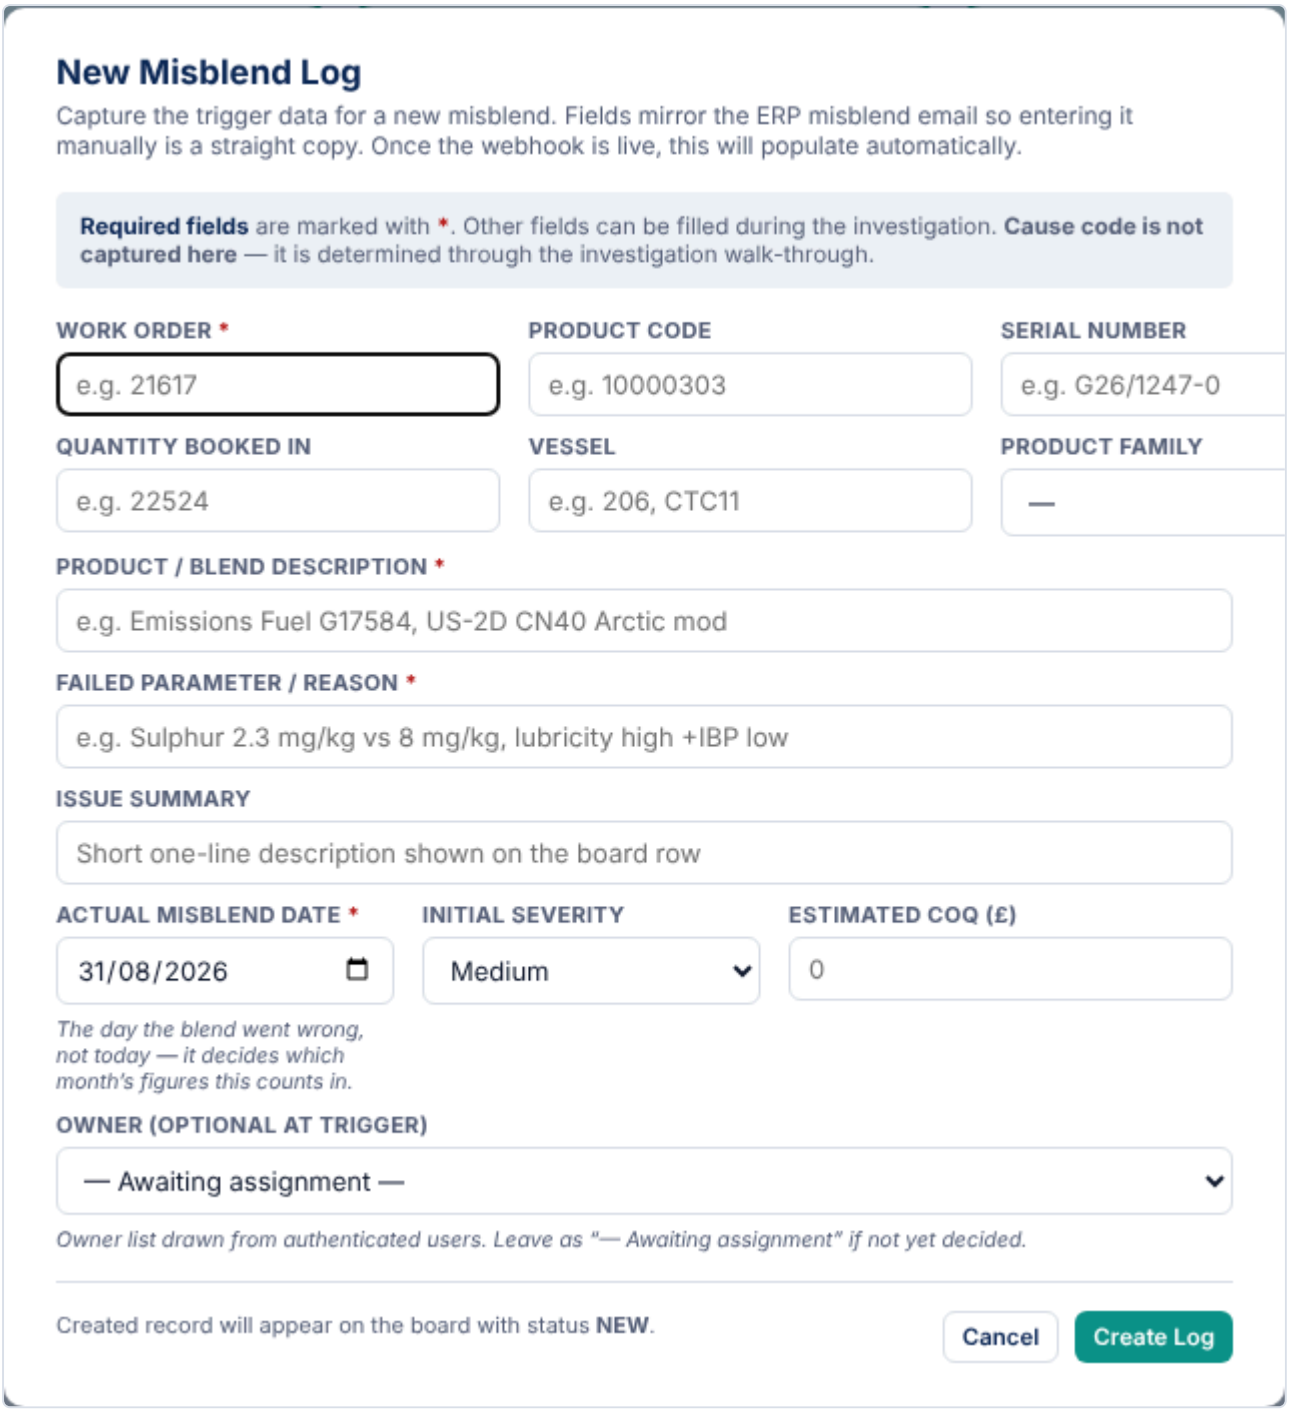
If a chart has more categories than fit, it caps the list and says in the axis title how many it left out. It never silently drops labels.

Underneath are two lists — **Live investigations** and **Recent outcomes** (the last seven days). Click any card to open the record. These are the same cards the wall board shows, built by the same code.

---

3. Raising a misblend record

+ New Record in the header. Records will arrive automatically once the ERP webhook is live; until then every one is raised here by hand.



The fields

FIELD	REQUIRED	NOTES
Work Order	Yes	If it already exists on another record the board asks whether you really want a second one
Product / Blend	Yes	
Failed Parameter	Yes	With the numbers — *Sulphur > 10 mg/kg*, *DVPE 65.2 vs 62.4*

FIELD	REQUIRED	NOTES
Actual misblend date	Yes	The day the blend went wrong. See below
Product Code	No	e.g. `10000303`
Serial Number	No	e.g. `G26/1247-0`
Quantity Booked In	No	
Vessel	No	e.g. `206`, `CTC11`
Product Family	No	From the admin-managed list
Owner	No	Defaults to *Awaiting assignment*
Severity	No	High / Medium / Low, defaults to Medium
Issue summary	No	
Cost of Quality	No	Better captured later, in the CoQ modal

Everything optional can be added from the record as you learn it. The point of raising it now is that the investigation clock starts and the record stops living in somebody's head.

The actual misblend date

This is **the day the blend went wrong**, not the day you are typing. It decides which month every figure counts the record in — the rate, the tile, the KPI row, the chart, the History date column.

Blends fail at month end and get written up in the first few days of the next one. A July misblend logged on 2 August belongs to July, against July's blends. Before this field existed it would have been August's, and August's rate would have carried a failure that was not its own.

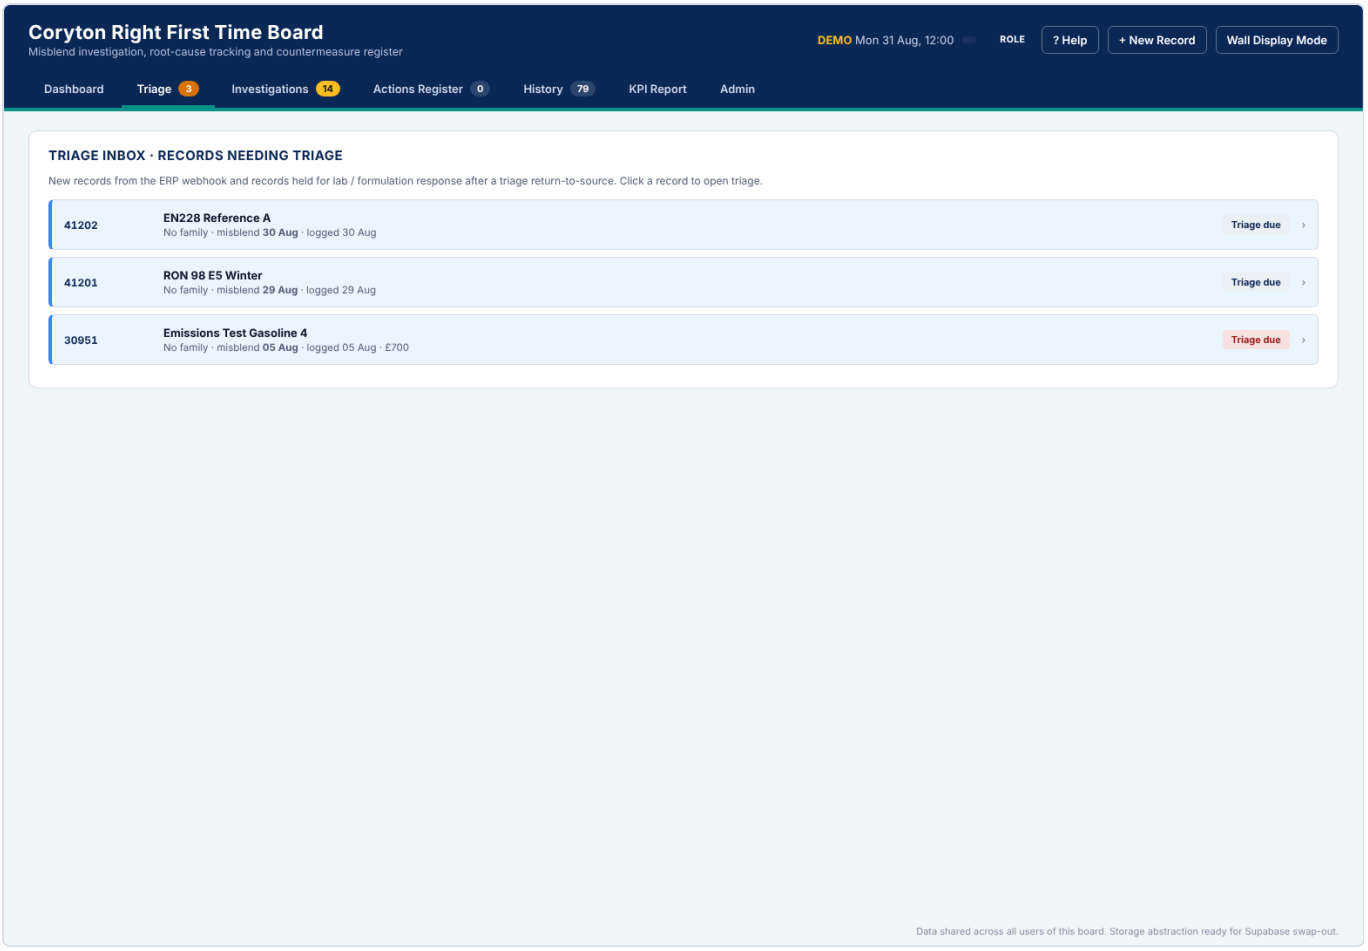
It defaults to today, which is right for most records, and it is shown again at the top of triage and in section 1 of the investigation so it can be corrected as soon as anyone notices.

The form **will not let you leave it blank and quietly call it today**. A blank silently becoming today is exactly the fault this field was raised to fix.

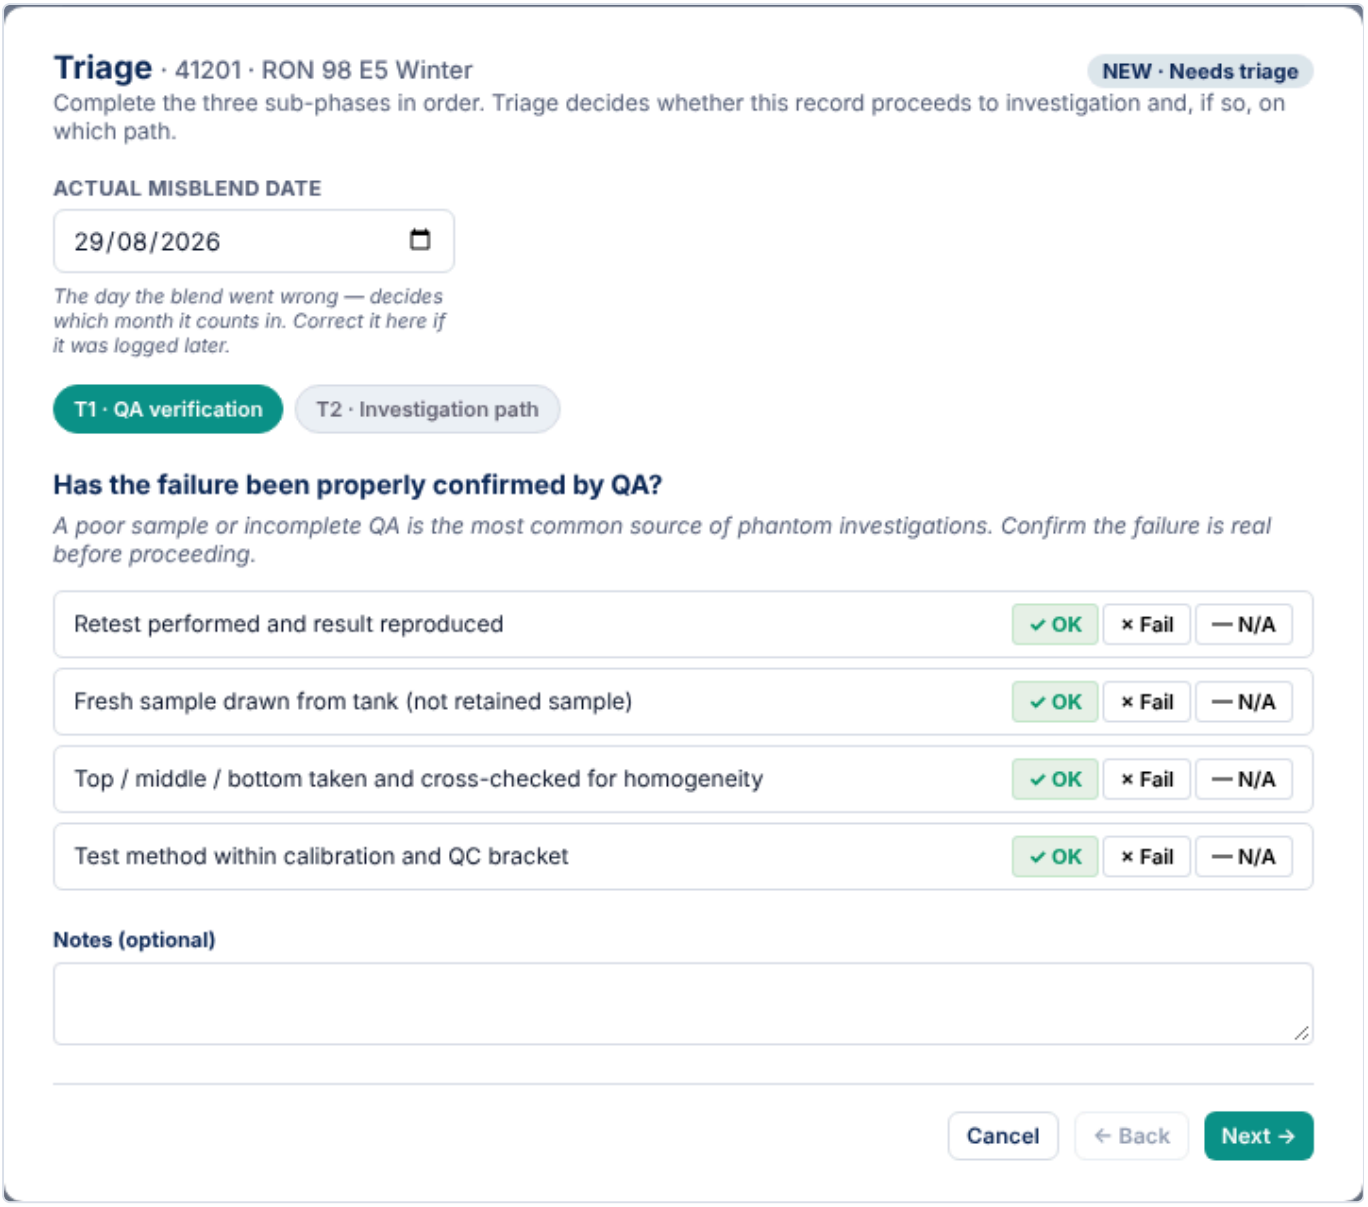
4. Triage

Triage exists because **the most common source of wasted investigation is a failure that was never real** — a poor sample, a tank that was not homogeneous, a test method out of bracket.

Records needing triage sit on the **Triage** tab, with a count on the tab itself. Click one to open triage.



T1 — QA verification



Four checks, each marked **OK**, **Fail** or **N/A**:

- Retest performed and result reproduced
- Fresh sample drawn from tank, not a retained sample
- Top / middle / bottom taken and cross-checked for homogeneity
- Test method within calibration and QC bracket

All confirmed and the record moves to path selection. Any gap and it goes back to the lab.

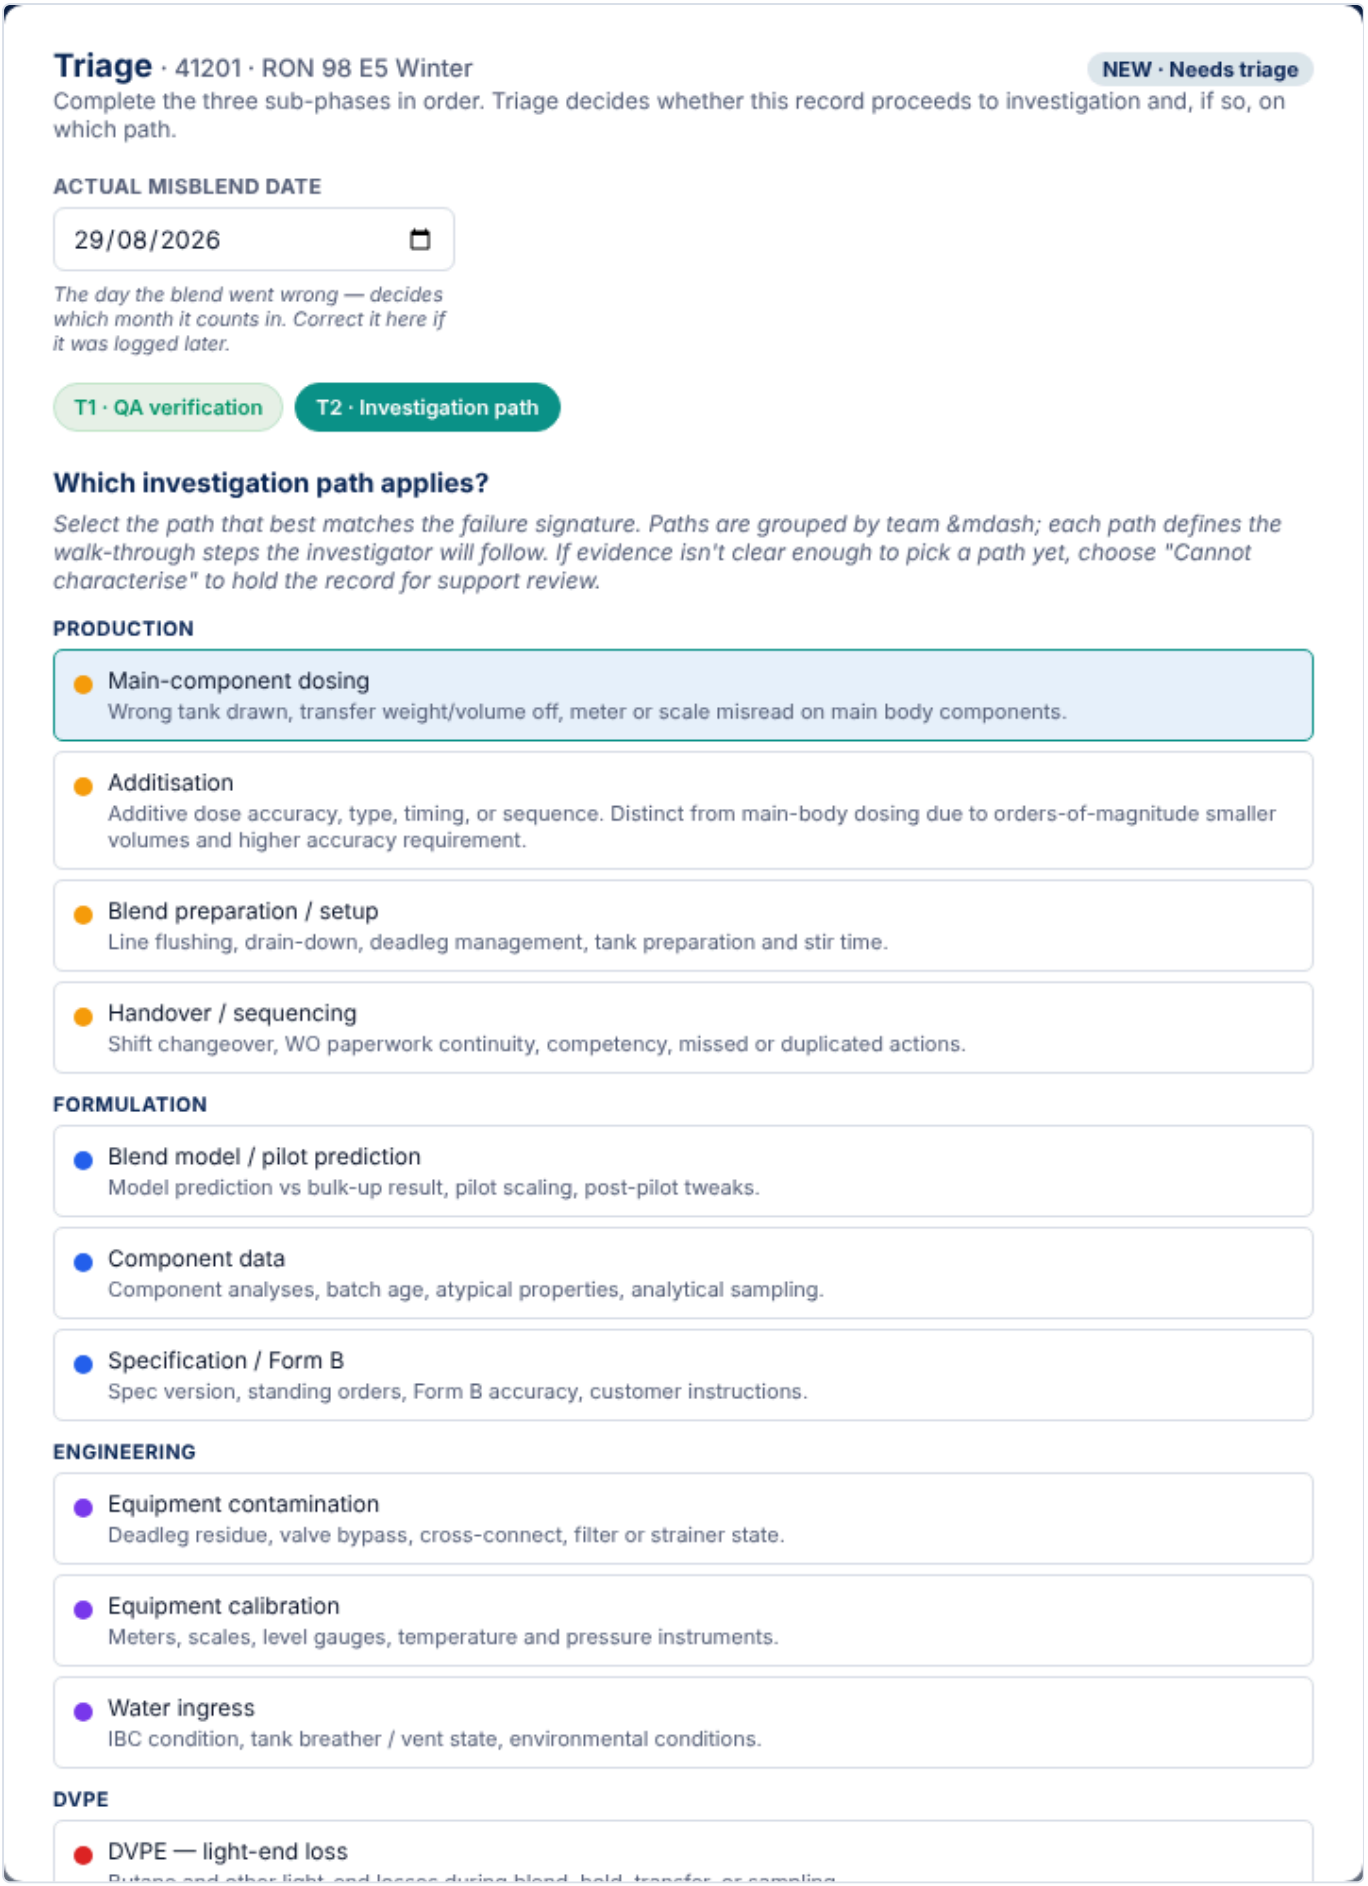
There is a notes box at each step, and the *Actual misblend date* sits at the top of the modal so it can be corrected here.

Returning it to the lab

Returning to source is **not a failure of the process — it is the process working**. A record held here has cost nothing beyond the triage itself.

The record stays on the Triage tab with its reason recorded, because it needs re-triaging once the lab has responded. Re-open triage from the same tab when they do; the pill at the top then reads **RE-TRIAGE** rather than *Needs triage*.

T2 — the investigation path



Pick the path that best matches the failure signature. There are fourteen, grouped by the team that usually owns them — but a path is a **failure mechanism, not a department**: *Main-component dosing*, *Equipment contamination*, *Light-end loss*.

You can select more than one. The walk-through then runs their steps as one sequence.

Each path defines the steps, the guidance, the evidence items and the cause codes those steps can conclude to. Administrators edit them in **Admin → Investigation Paths**, and changes flow straight through to triage.

When you cannot tell yet

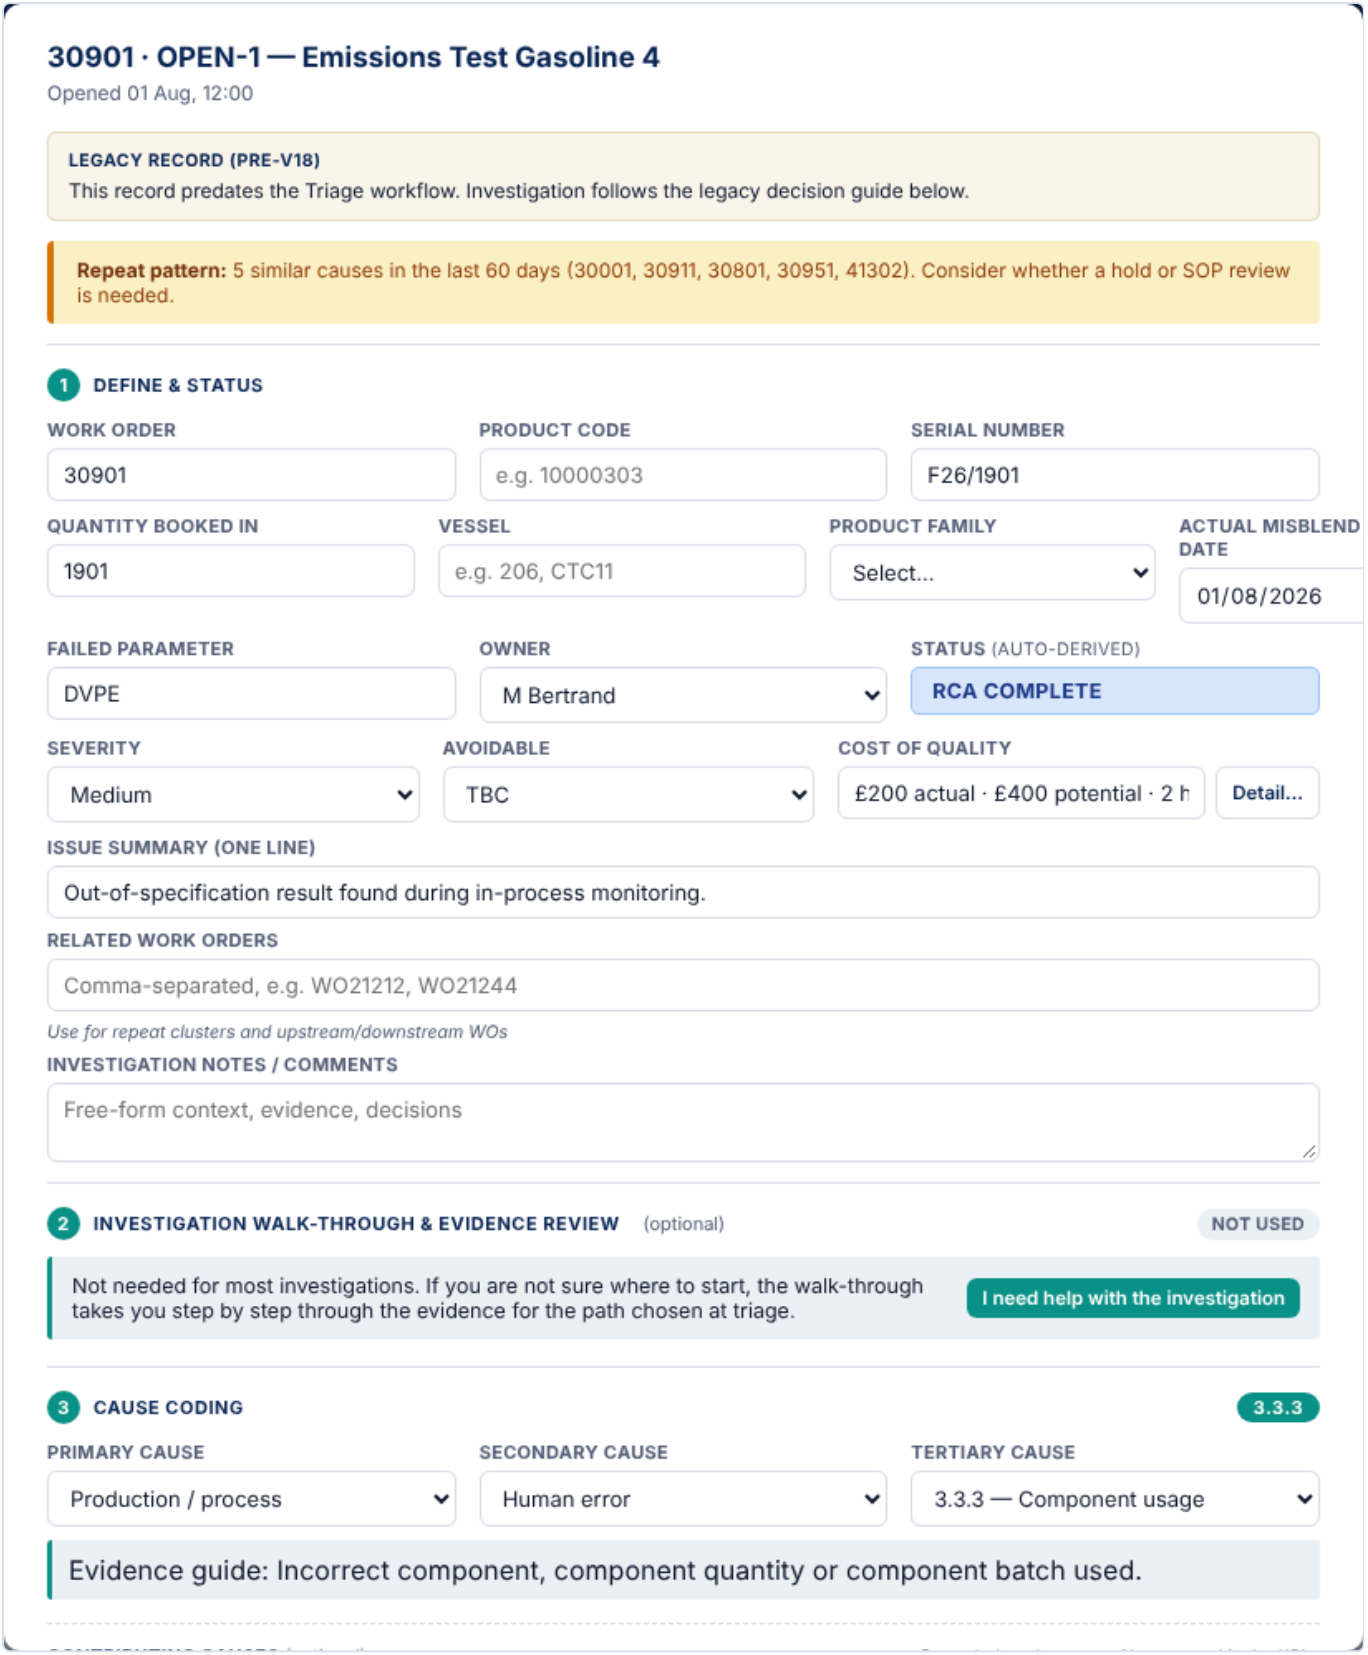
Choose **Cannot characterise from current evidence — hold for support** and the record is held for support review rather than sent down a path on a guess. It stays on the Triage tab until somebody can characterise it.

This is a better answer than picking the most likely path and hoping. A wrong path costs the investigator's time and leaves a walk-through in the record that has to be explained later.

---

5. Running the investigation

Open any record and it is the same seven numbered sections, in the order the work happens.



#	SECTION	WHAT GOES IN IT
1	Define & Status	The ERP identifiers, failed parameter, owner, severity, misblend date
2	Investigation Walk-through & Evidence Review	Optional — the guided steps for the chosen path, if you ask for them

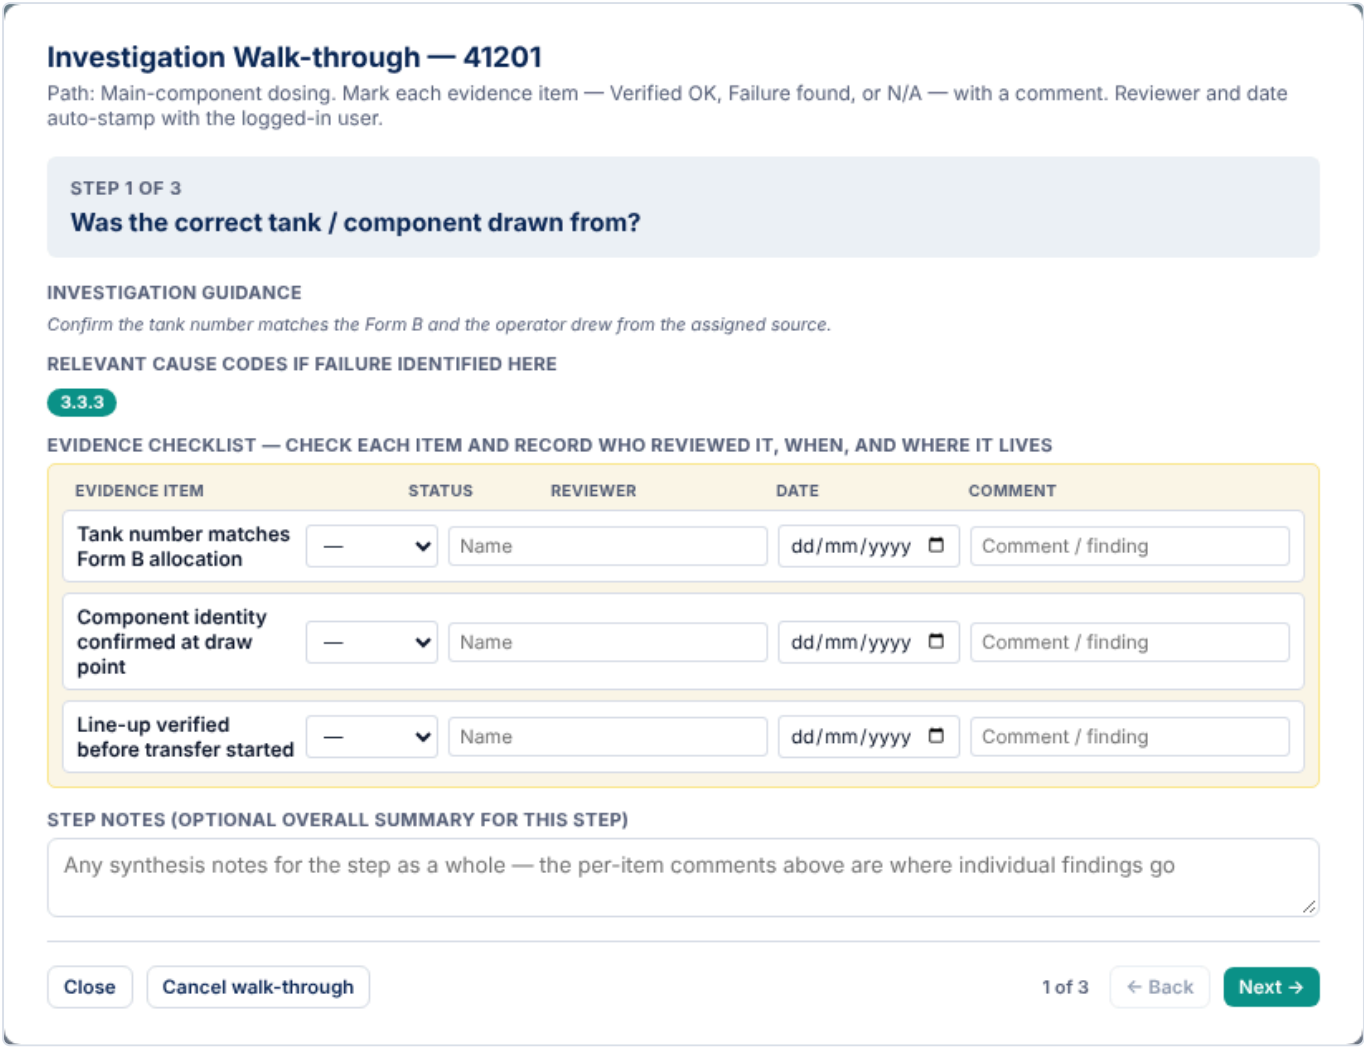
#	SECTION	WHAT GOES IN IT
3	Cause coding	The taxonomy code, plus contributing causes
4	5 Whys	Symptom to the thing that actually needs fixing
5	Containment	What was done with the affected product straight away
6	Actions	Corrective and preventive countermeasures
7	Root Cause & Sign-off	The narrative, and the two-person approval

Status is derived, never typed

The status field has no dropdown. It is computed from how far the record has actually got — triaged, walked, coded, prepared, approved — so it can never claim a record is further along than it is.

To move a record on, do the next piece of work. The status follows.

The walk-through



The walk-through is optional. Most investigations do not need it, so section 2 starts closed. If you are not sure where to start, press **I need help with the investigation** and it opens: one step at a time, each with a question, guidance, and the evidence items that settle it. Mark each item, answer the step, add notes, move on.

Once you have answered anything in it, the walk-through has to be finished before the report can be signed off — a half-walked path in a report reads as work left undone. If you no longer need it, tick **Stop using the walk-through**, or press **Cancel walk-through** — it is beside *Close* inside the walk-through, and in section 2 once you have

started. The steps you answered stay on the record, your name and the time are recorded against the decision, and the walk-through no longer holds up sign-off. The board asks first, with two buttons — **Keep the walk-through** and **Cancel walk-through** — so there is no guessing what *OK* means. Untick the box to pick it up again. Cancelling before you have answered anything just closes it, as if you had never asked. A report that never used it prints *Not used* in that section.

Each step stores **a snapshot of its own wording** as you walk it. If an administrator rewrites that step next year, your report still renders the question you were actually asked, not the current one.

When the path was wrong

If you walk a path and the root cause is not there, use **Try another investigation path**. The attempt you were on is marked superseded and a new one starts.

Nothing is thrown away. The earlier attempt's steps, evidence and notes stay on the record and are visible in the attempt history — because an investigation that ruled two things out before finding the third is a stronger record than one that went straight there.

Two people, one record

A record is locked while somebody is editing it. Open one that is already held and you get a yellow banner naming who has it; you can read everything but not change it, and their saves appear to you within a few seconds.

A lock goes stale after **five minutes** without activity, after which anyone can take the record. That is the backstop for a tab left open over lunch.

The lock is a courtesy against two people typing over each other, not a permission. Any investigator can edit any record — there is no per-record ownership.

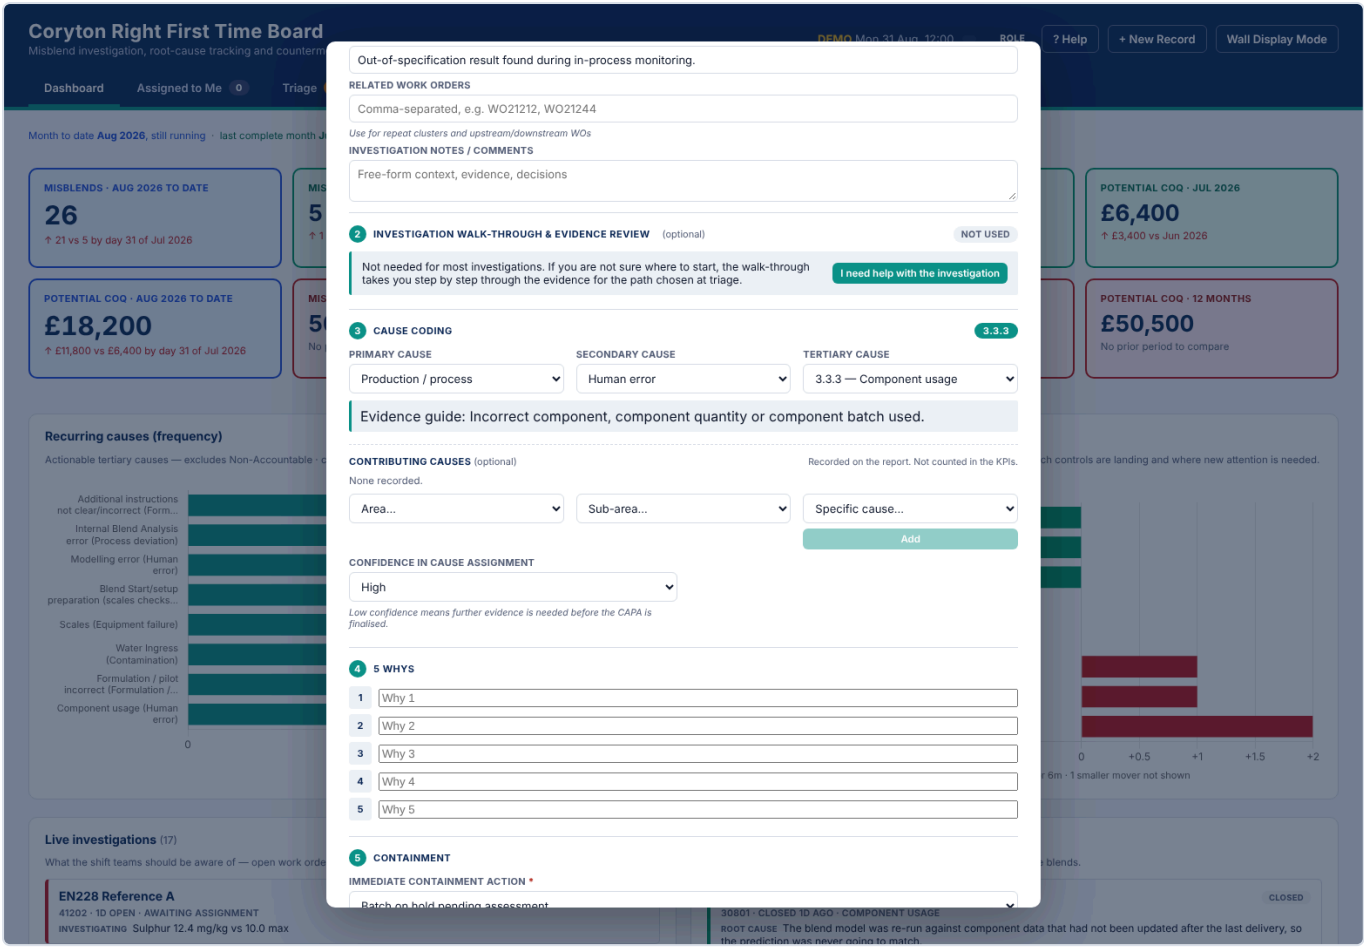
A tab left open does not keep up by itself. The board reads the records when it loads and does not read them again, so a tab open since the morning is showing the morning. Two things stop that costing anyone their work:

- **Opening a record reads it again.** If somebody saved it since your page loaded, you see their version, and a message says so.
- **A save never lands on top of work you have not seen.** If the record has moved on since you loaded it, the save is refused and a red bar names the record — see *A red bar says someone else saved a newer version* below.

Reload the page to see everything as it is now.

---

6. Cause coding and 5 Whys



The taxonomy

Causes are coded **Area › Sub-area › Specific**, shown as a numeric code such as `3.2.4`. Choosing the three dropdowns in section 3 sets the code; if you used the walk-through, it will usually have proposed one already.

Administrators maintain the list in **Admin → Cause Catalogue**, with an evidence guide against each code. Historic records keep the code they were given even if it is renamed later.

Primary and contributing causes

One record has **one primary cause** — the thing that, fixed, would have prevented it. Anything else that made it more likely or worse goes in as a **contributing cause**.

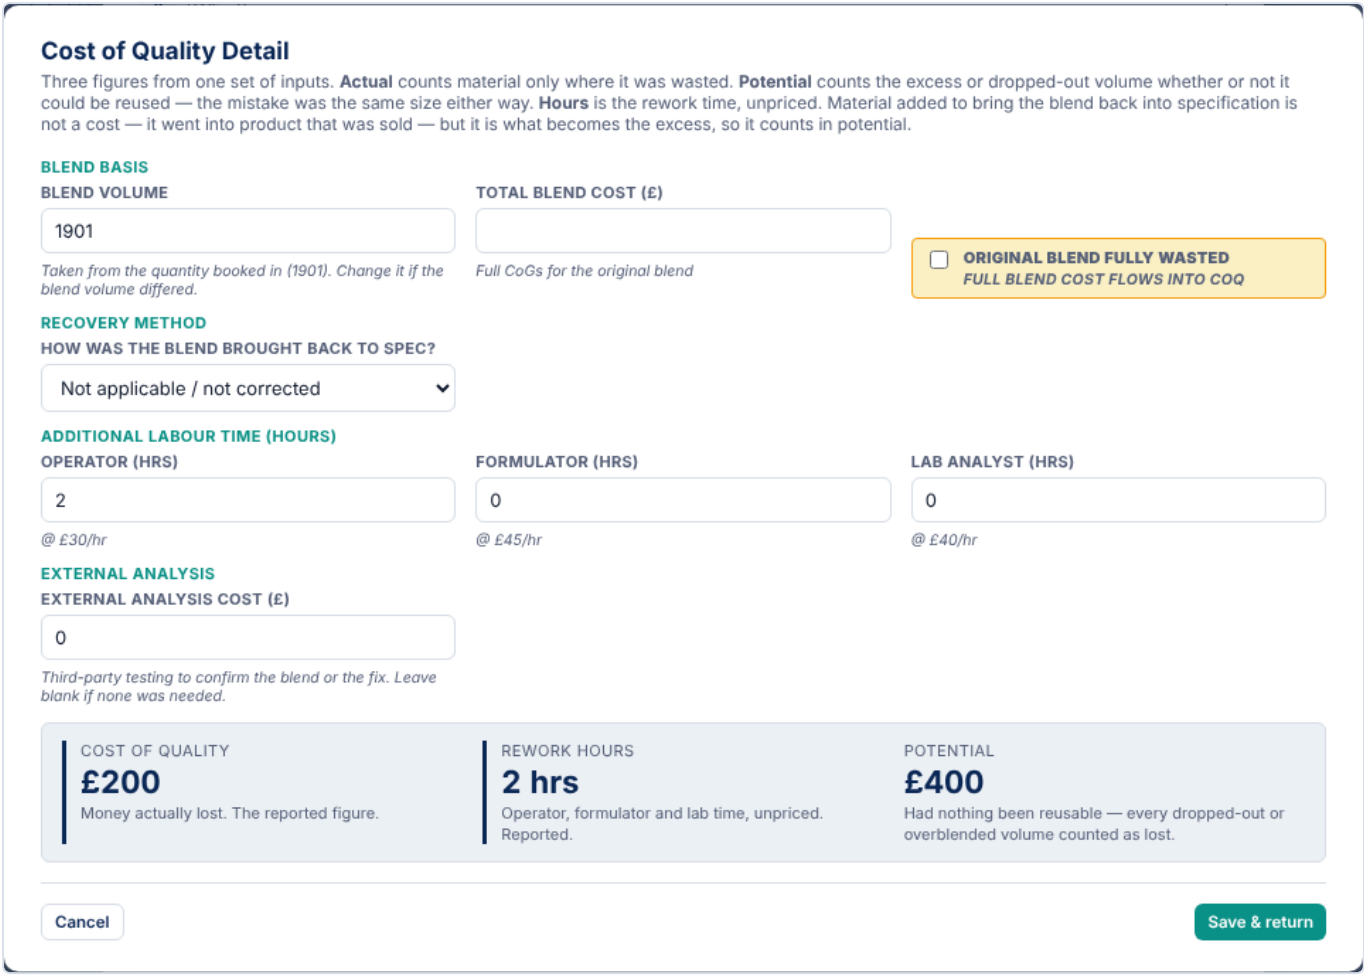
Only the primary code drives the KPI. Contributing causes are there so the record tells the truth about a failure that had more than one ingredient.

5 Whys

Five boxes: the symptom at the top, and each answer becomes the next question. Stop when you reach something you can actually change — a procedure, a check, a piece of equipment, a piece of data.

If the fifth why is still a person making a mistake, the chain has stopped too early. *"Operator misread the meter"* is a symptom; *"the meter reads in a unit the Form B does not use"* is something you can fix.

7. Cost of quality



Three figures, from one set of inputs

FIGURE	WHAT IT COUNTS	WHERE IT APPEARS
Actual	Money genuinely lost — material that was *wasted*, external analysis, rework labour	KPI tab, History, the record, the report
Rework hours	Operator, formulator and analyst time, unpriced	Its own tile, the KPI cost table, the record, the report
Potential	The same as actual, but material counts *whether or not it was wasted*. Never less than actual	**The only cost on the tiles** — the dashboard and both wall boards — plus the KPI cost table

Why the tiles show potential

Reuse is luck, not process. A misblend whose excess happened to be reusable this time is the same process failure as one whose excess went to slops; the difference is whether there was somewhere to put it.

The actual cost sat on the dashboard beside potential for one build and read exactly as *"we messed up but it was all fine"*, so it came off. Actual is still there in full on the KPI tab, where you are looking at the arithmetic rather than at a headline.

Wasted, reused, or not known

Whether the excess was wasted or reused is the single input that moves the figures most, and three states are possible:

- **Wasted** — counts in actual and in potential.
- **Reused** — counts in potential only.
- **Not known** — a number of imported historic records carry value at risk with no flag either way.

An unknown record is counted at the figure the original register carried and **flagged as unconfirmed** rather than guessed in either direction. The KPI cost table has an *Unconfirmed* column, History has a *Cost outcome* filter, and the modal lets you settle one from the register's own components when you know the answer.

Material added to bring a blend back into specification

There is deliberately no input for it, and it is not in actual: that material went into product that was sold.

But it *is* the excess — it is what would have been lost had the excess not been reusable — so it is exactly what potential adds back. That is the whole difference between the two figures.

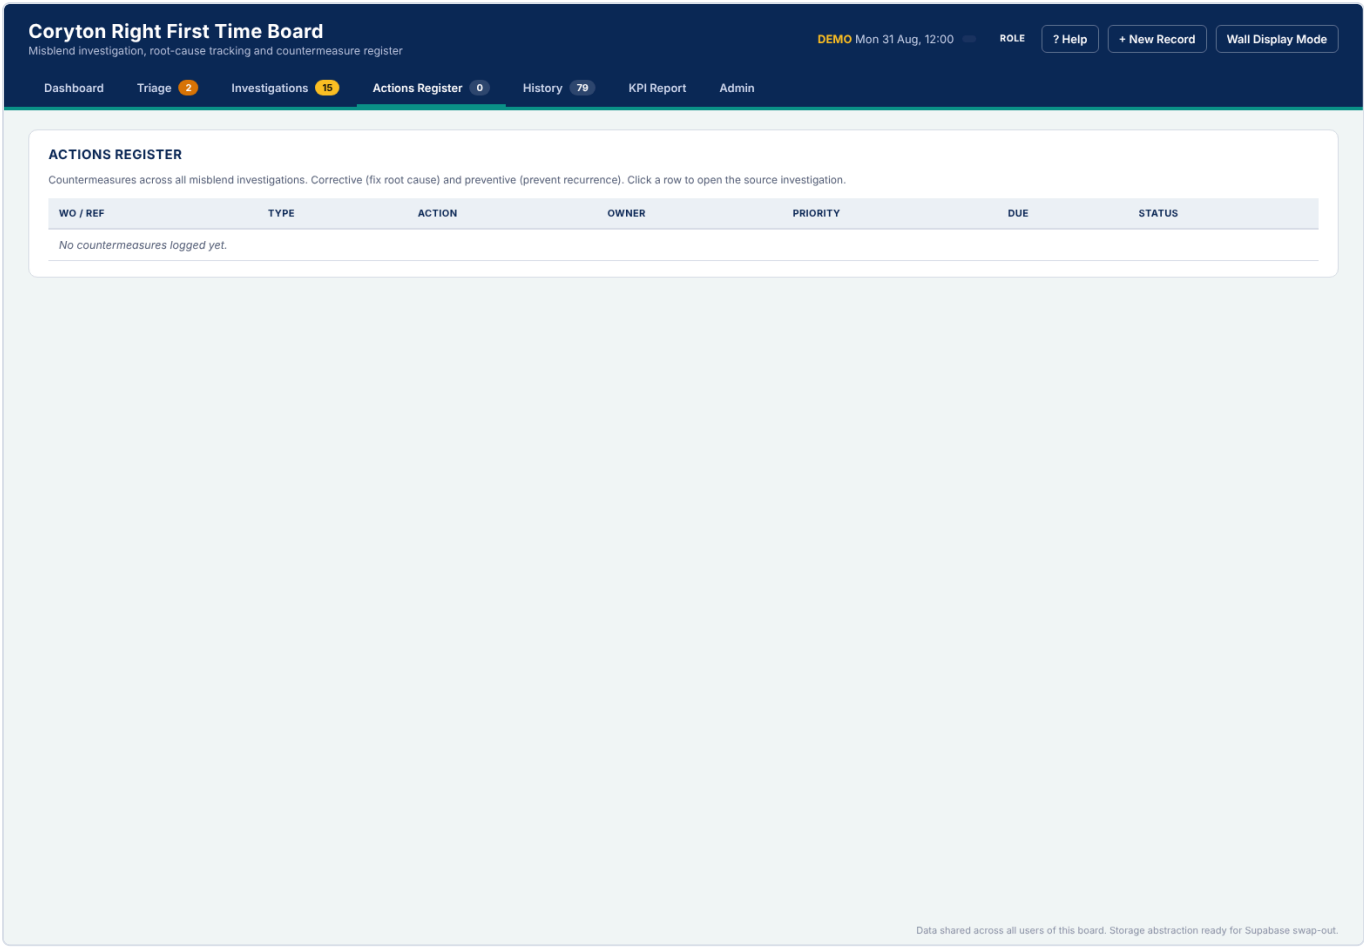
Checking a month's figure

KPI → Cost of Quality → the records behind a month takes a month apart record by record: value at risk, lost, analysis, labour, hours, and where each came from. The total must equal that month's row in the table above it.

Nothing on that table is recomputed. The cost and potential shown are the stored figures, with the parts listed so the sum can be followed.

---

8. The actions register

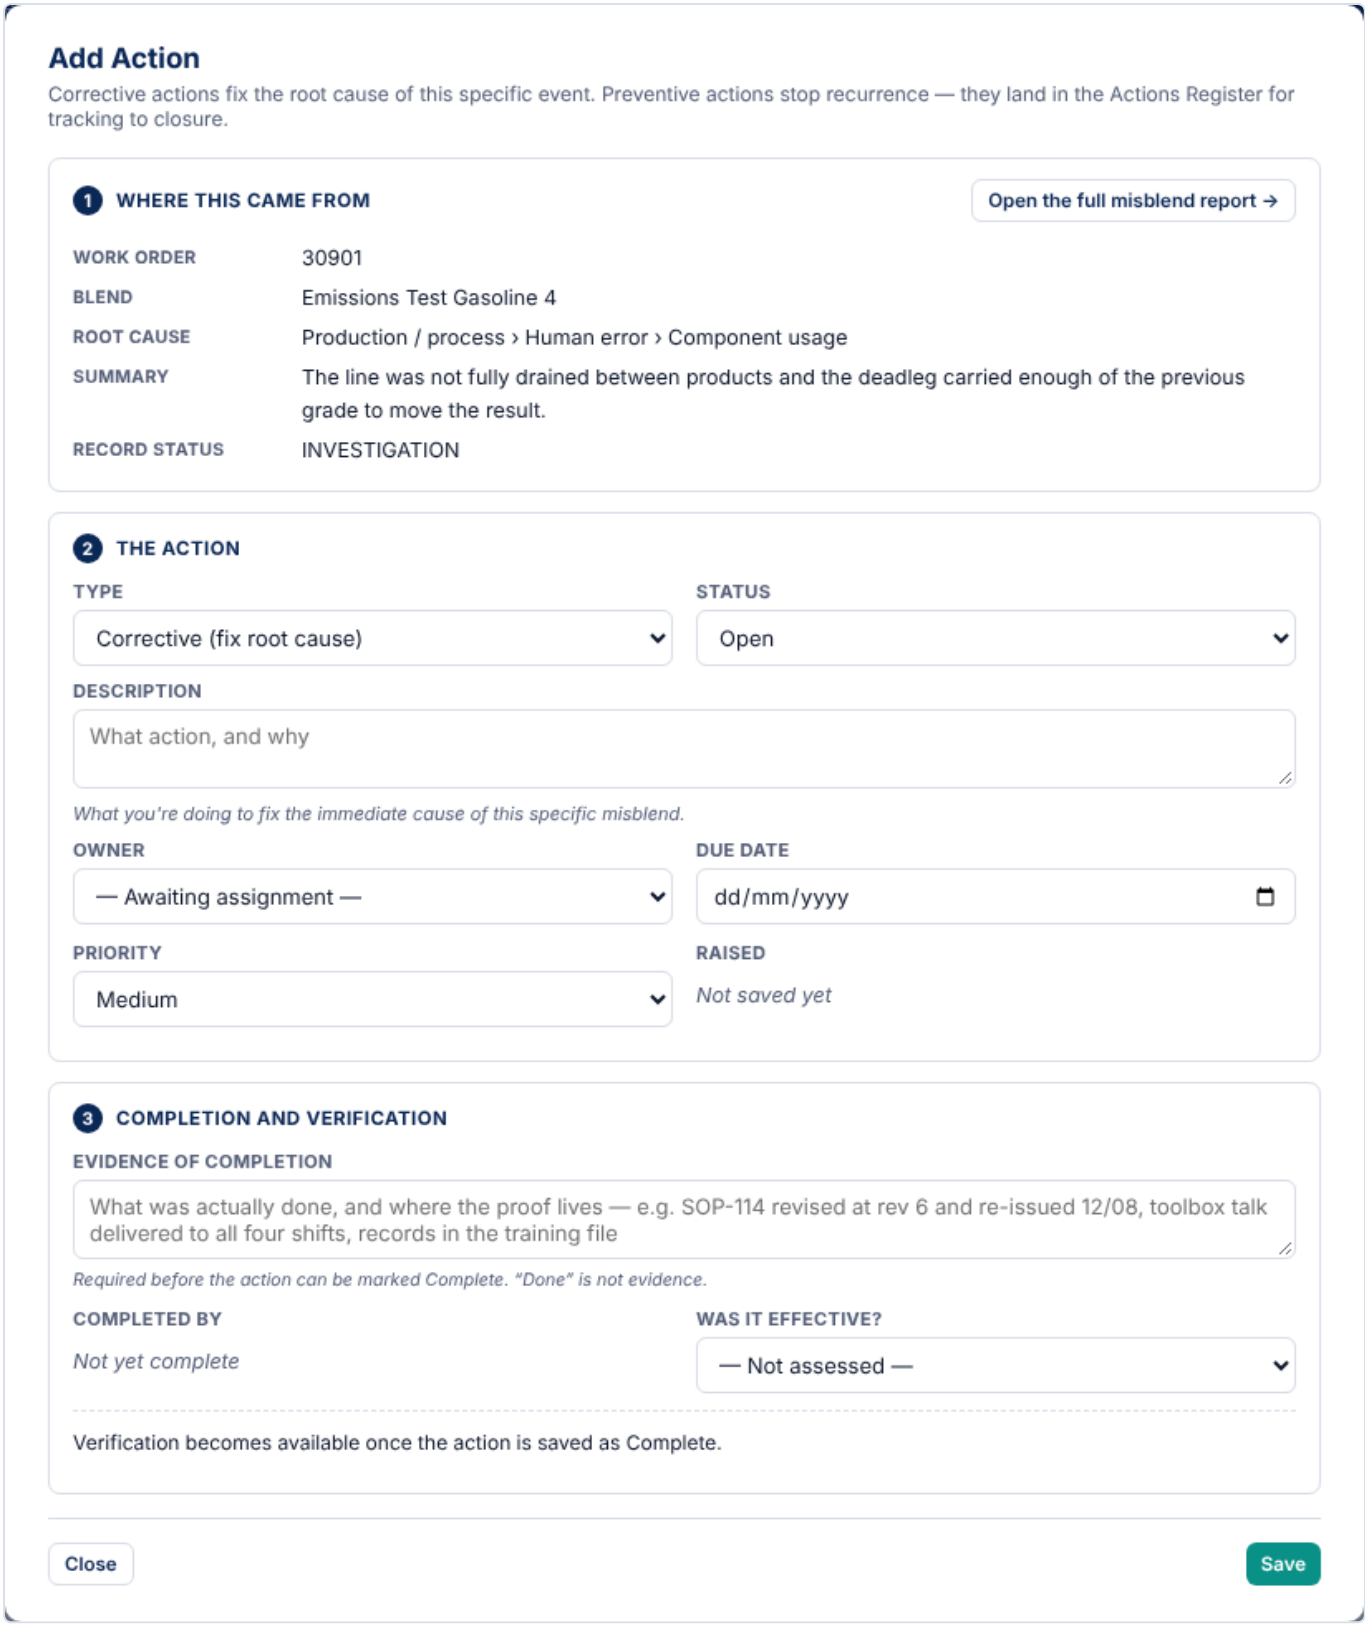


Corrective and preventive

Corrective fixes the root cause of this failure. **Preventive** stops it recurring, usually somewhere broader than the one blend that triggered it.

Most investigations need one of each. An investigation that raises only a corrective action has fixed today and left tomorrow alone.

Raising and verifying one



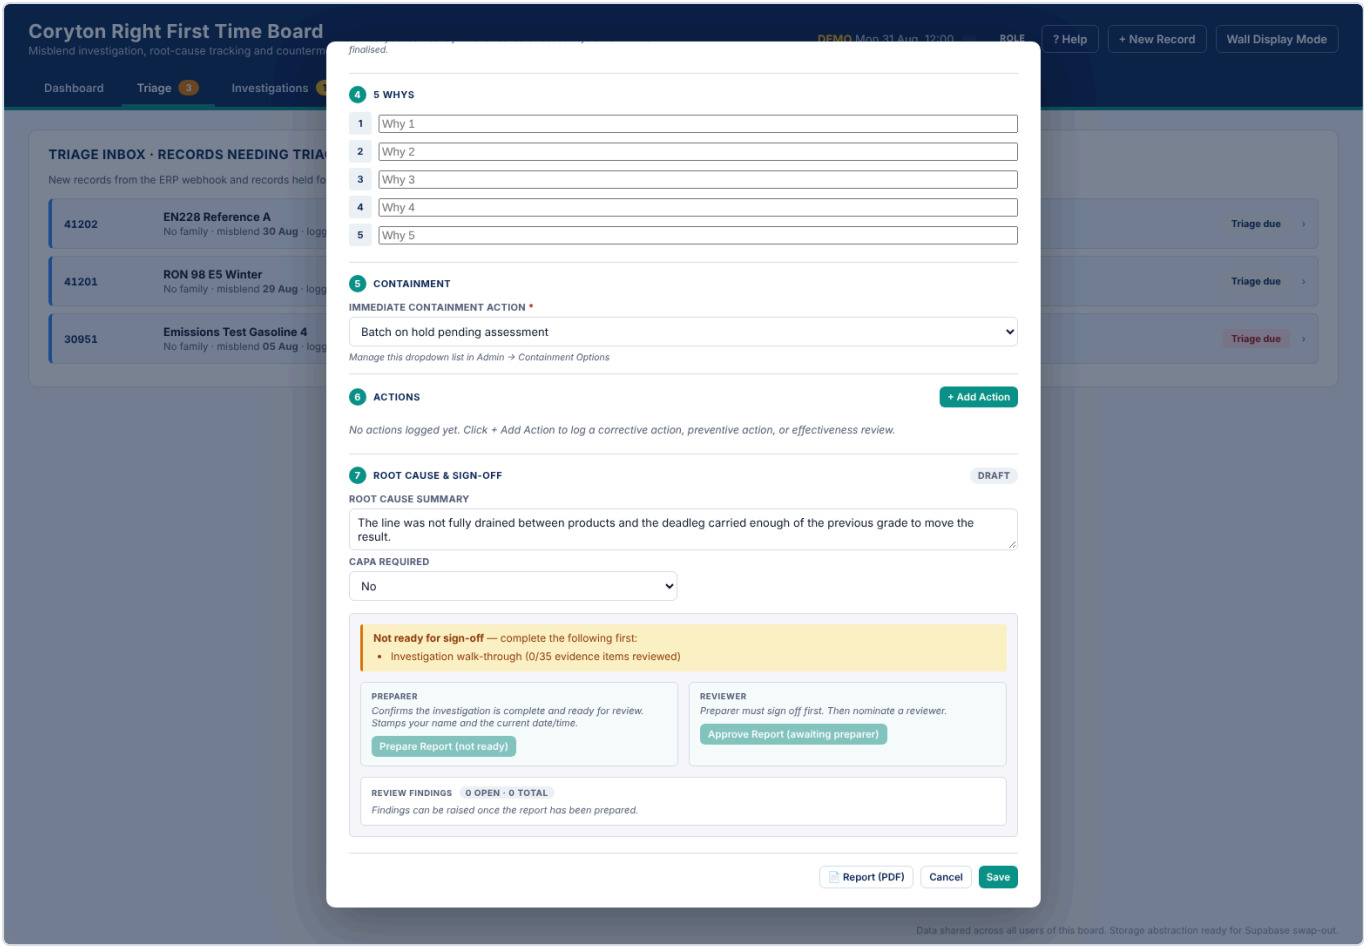
Actions are raised from section 6 of a record. Each carries an owner, a priority, a due date and the text of what is to be done.

An action is not finished when it is done — it is finished when somebody has confirmed it was done **and that it worked**. The action carries a verification block with who verified it and when. Closing one without verification leaves the register saying a countermeasure is in place when nobody has checked that it is.

The **Actions Register** tab shows every action from every investigation, with a count of what is outstanding on the tab itself, so a due date is visible without opening the record it belongs to. Click a row to open the source

investigation.

9. Sign-off and review



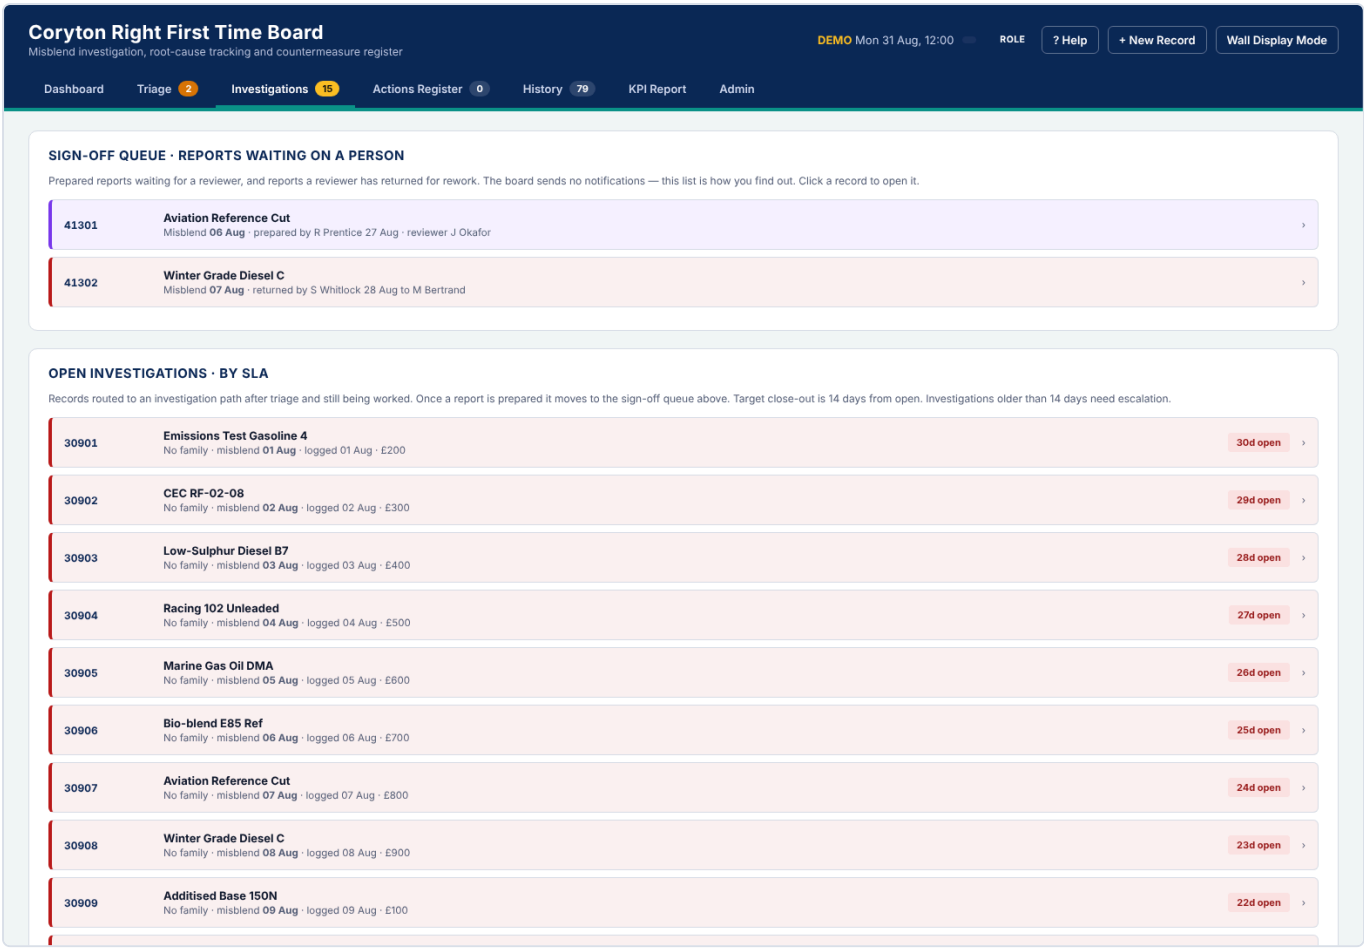
The two-person rule

A report is **prepared** by one person and **approved** by another. The board disables Approve for whoever prepared it, whatever their role.

The preparer nominates a reviewer. Only the preparer, the record's owner or an administrator can make or change that nomination — a nomination anyone could change would decide nothing.

The rule is enforced in the browser, not in the database. For an internal tool among colleagues that is a reasonable position; it is not audit-proof, and it is written down as such in `docs/SECURITY.md`.

The sign-off queue



Nominating a reviewer notifies nobody. There is no email. The sign-off queue at the top of the Investigations tab is how a reviewer finds out a report is waiting, and how a preparer finds out one came back — the ones waiting on you are flagged and listed first.

Telling your reviewer is still a manual step. The board says so rather than pretending otherwise.

Findings — question, challenge, issue

Approving is not the reviewer's only move. Against a prepared report the nominated reviewer, or an administrator, can raise a **finding** pointed at a section: a *question*, a *challenge* to the narrative, or an *issue* with the report.

The preparer answers it on the record. The reviewer resolves it. Only the person who raised a finding can withdraw it.

Approve is disabled while any finding is open, and an answered finding still counts as open until the reviewer resolves it. A challenge that could be signed past would be decoration.

With nobody nominated, only an administrator can approve — which is the nudge to nominate one.

Returning a report for rework

If the report needs real work rather than an answer, the reviewer **returns it** with a reason. The prepare stamp is cleared, the return is recorded with who, when, why and how many findings were open, and the preparer prepares again as the next round.

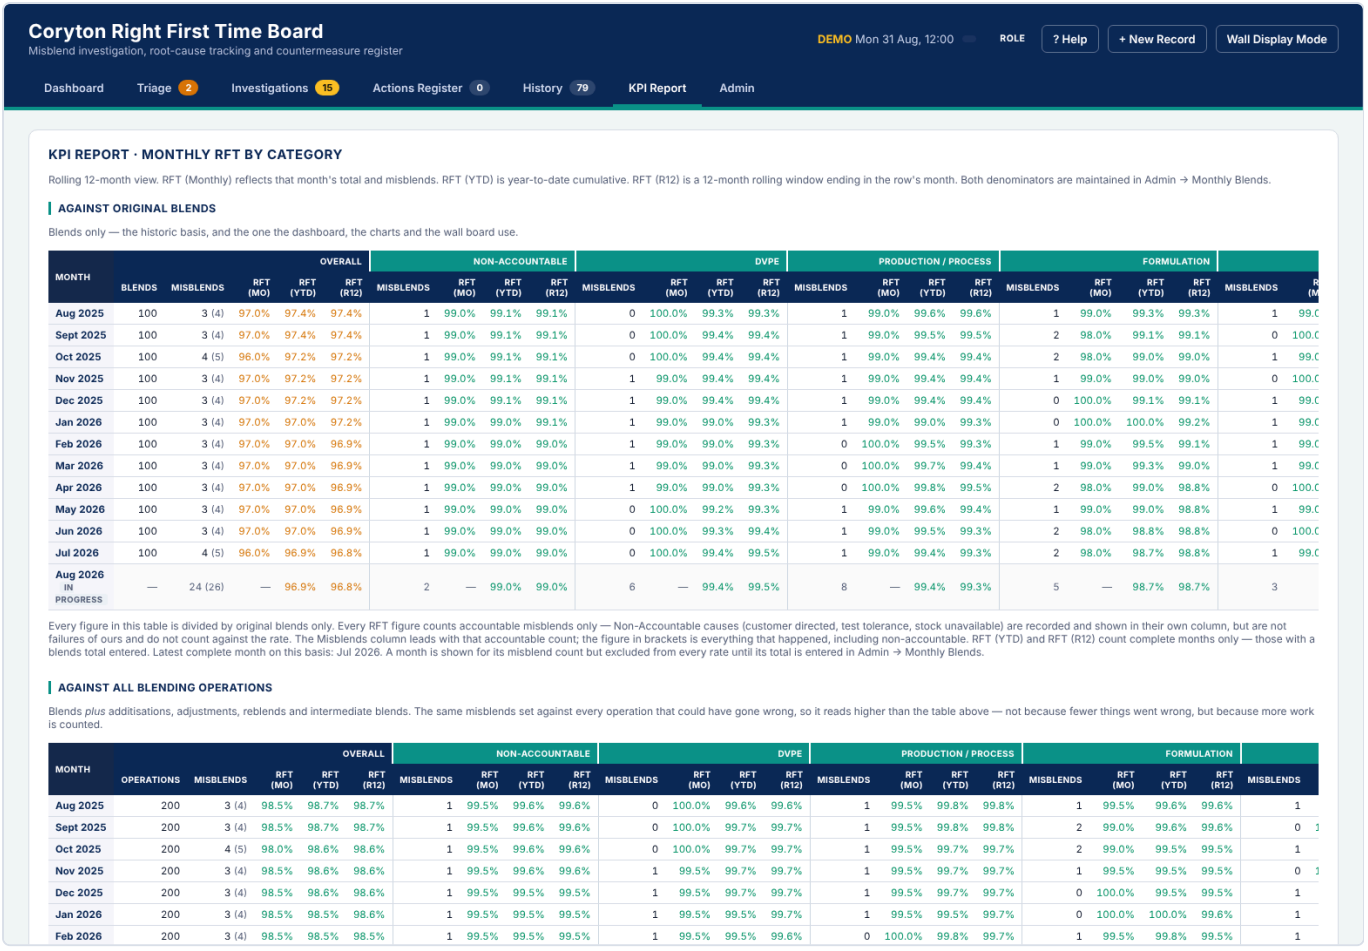
Every finding records the round it was raised in, so a report returned twice reads as two rounds rather than one long argument.

Nothing is deleted

Resolved and withdrawn findings stay on the record and print in the report's **Review Record**. That is what turns two names on a page into evidence that a review actually happened.

---

10. The KPI report



What RFT % is

```
RFT % = (blends in month − accountable misblends in month) ÷ blends in month
```

The numerator counts **accountable** misblends only. The denominator is the blend count for that month, typed in from the ERP monthly report by an administrator.

Complete months only

Blend counts arrive at month end. **A month with no blend count entered does not appear in the KPI report at all** — it is excluded rather than shown as zero.

That is correct, and it is easy to mistake for missing data. A month shown as zero would divide a complete numerator by a missing denominator and every rate would read low.

If a month you expect is not there, the answer is almost always that nobody has entered its blend count yet.

Two denominators

A misblend is a misblend whichever operation produced it, so the numerator is the same either way. What changes is what it is divided by:

BASIS	COUNTS
Original blends	Original blends only. The historic basis, and the default everywhere else on the board
All blending operations	Blends plus additisations, adjustments, reblends and intermediate blends

The KPI tab shows both, as two copies of the same table one under the other. Every tile, chart and note elsewhere — and the wall board — uses original blends.

A month absent from one series does not borrow the other. A month with blends but no operations total has no operations rate at all, so a month can be complete on one basis and still in progress on the other. Filling it in from the blend count would claim there were no additisations that month, which is the one thing we know we do not know.

Reading the Misblends column

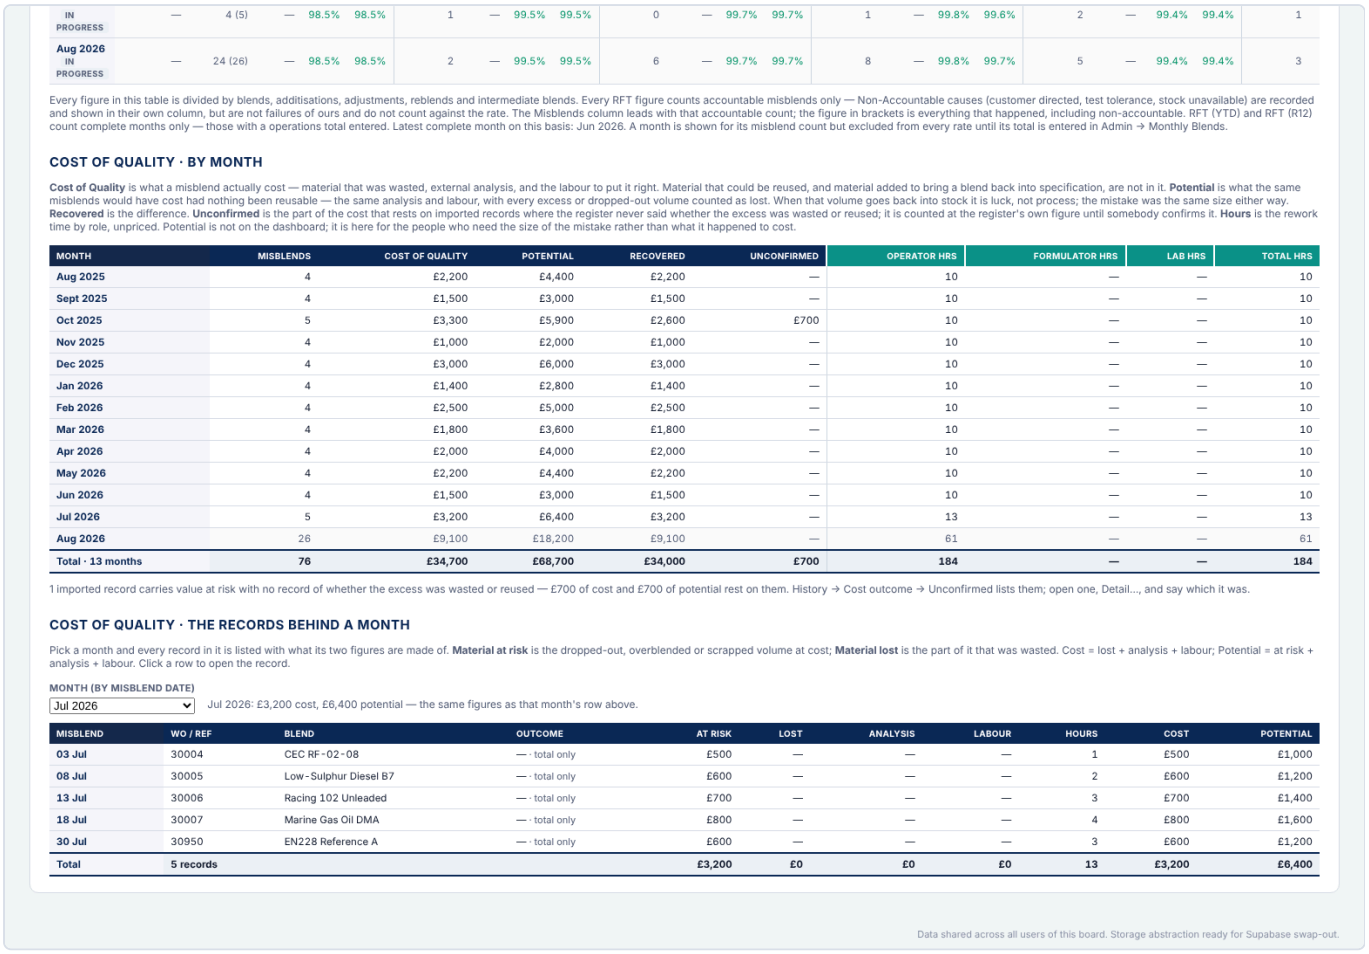
It reads `3 (4)` : the **accountable** count first, the total including non-accountable causes in brackets.

The number doing the arithmetic goes first. It was the other way round until a user pointed out that the column and the rate beside it looked like they disagreed.

The three windows

Overall and every category carry **RFT (mo)**, **(YTD)** and **(R12)**. All three count complete months only. Overall had just the month until a user pointed out that the one figure anybody actually reports was the one nobody could read off the table.

The cost tables



IN PROGRESS	—	4 (5)	—	98.5%	98.5%	1	—	99.5%	99.5%	0	—	99.7%	99.7%	1	—	99.8%	99.6%	2	—	99.4%	99.4%	1
Aug 2026 IN PROGRESS	—	24 (26)	—	98.5%	98.5%	2	—	99.5%	99.5%	6	—	99.7%	99.7%	8	—	99.8%	99.7%	5	—	99.4%	99.4%	3

Every figure in this table is divided by blends, additisations, adjustments, reblends and intermediate blends. Every RFT figure counts accountable misblends only — Non-Accountable causes (customer directed, test tolerance, stock unavailable) are recorded and shown in their own column, but are not failures of ours and do not count against the rate. The Misblends column leads with that accountable count; the figure in brackets is everything that happened, including non-accountable. RFT (YTD) and RFT (R12) count complete months only — those with a operations total entered. Latest complete month on this basis: Jun 2026. A month is shown for its misblend count but excluded from every rate until its total is entered in Admin → Monthly Blends.

COST OF QUALITY · BY MONTH

Cost of Quality is what a misblend actually cost — material that was wasted, external analysis, and the labour to put it right. Material that could be reused, and material added to bring a blend back into specification, are not in it. **Potential** is what the same misblends would have cost had nothing been reusable — the same analysis and labour, with every excess or dropped-out volume counted as lost. When that volume goes back into stock it is luck, not process; the mistake was the same size either way. **Recovered** is the difference. **Unconfirmed** is the part of the cost that rests on imported records where the register never said whether the excess was wasted or reused; it is counted at the register's own figure until somebody confirms it. **Hours** is the rework time by role, unpriced. Potential is not on the dashboard; it is here for the people who need the size of the mistake rather than what it happened to cost.

MONTH	MISBLENDS	COST OF QUALITY	POTENTIAL	RECOVERED	UNCONFIRMED	OPERATOR HRS	FORMULATOR HRS	LAB HRS	TOTAL HRS
Aug 2025	4	£2,200	£4,400	£2,200	—	10	—	—	10
Sept 2025	4	£1,500	£3,000	£1,500	—	10	—	—	10
Oct 2025	5	£3,300	£5,900	£2,600	£700	10	—	—	10
Nov 2025	4	£1,000	£2,000	£1,000	—	10	—	—	10
Dec 2025	4	£3,000	£6,000	£3,000	—	10	—	—	10
Jan 2026	4	£1,400	£2,800	£1,400	—	10	—	—	10
Feb 2026	4	£2,500	£5,000	£2,500	—	10	—	—	10
Mar 2026	4	£1,800	£3,600	£1,800	—	10	—	—	10
Apr 2026	4	£2,000	£4,000	£2,000	—	10	—	—	10
May 2026	4	£2,200	£4,400	£2,200	—	10	—	—	10
Jun 2026	4	£1,500	£3,000	£1,500	—	10	—	—	10
Jul 2026	5	£3,200	£6,400	£3,200	—	13	—	—	13
Aug 2026	26	£9,100	£18,200	£9,100	—	61	—	—	61
Total · 13 months	**76**	**£34,700**	**£68,700**	**£34,000**	**£700**	**184**	**—**	**—**	**184**

1 imported record carries value at risk with no record of whether the excess was wasted or reused — £700 of cost and £700 of potential rest on them. History → Cost outcome → Unconfirmed lists them; open one, Detail…, and say which it was.

COST OF QUALITY · THE RECORDS BEHIND A MONTH

Pick a month and every record in it is listed with what its two figures are made of. **Material at risk** is the dropped-out, overblended or scrapped volume at cost; **Material lost** is the part of it that was wasted. Cost = lost + analysis + labour; Potential = at risk + analysis + labour. Click a row to open the record.

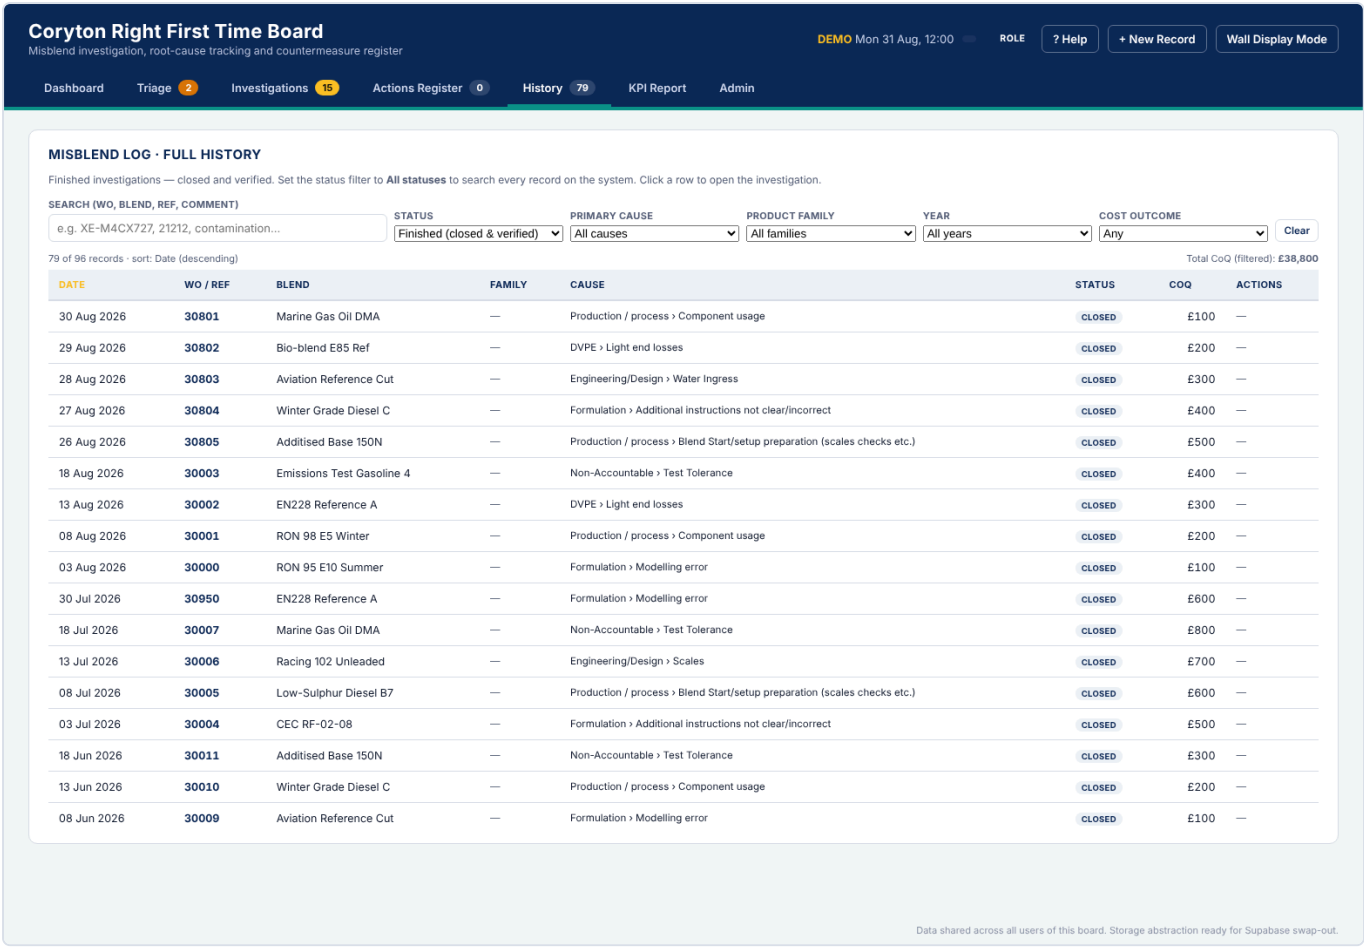
MONTH (BY MISBLEND DATE) Jul 2026 — Jul 2026: £3,200 cost, £6,400 potential — the same figures as that month's row above.

MISBLEND	WO / REF	BLEND	OUTCOME	AT RISK	LOST	ANALYSIS	LABOUR	HOURS	COST	POTENTIAL
03 Jul	30004	CEC RF-02-08	— · total only	£500	—	—	—	1	£500	£1,000
08 Jul	30005	Low-Sulphur Diesel B7	— · total only	£600	—	—	—	2	£600	£1,200
13 Jul	30006	Racing 102 Unleaded	— · total only	£700	—	—	—	3	£700	£1,400
18 Jul	30007	Marine Gas Oil DMA	— · total only	£800	—	—	—	4	£800	£1,600
30 Jul	30950	EN228 Reference A	— · total only	£600	—	—	—	3	£600	£1,200
Total	**5 records**			**£3,200**	**£0**	**£0**	**£0**	**13**	**£3,200**	**£6,400**

Data shared across all users of this board. Storage abstraction ready for Supabase swap-out.

Beneath the rate tables: cost of quality by month, with **Potential**, **Recovered** and **Unconfirmed** columns, and underneath it *the records behind a month* — see §7.

11. History and reports



Finding a record

The **History** tab is every record the board holds, not just the closed ones. Filter by free-text search, status, primary cause, product family, year and cost outcome, and clear the lot with one button.

It opens with the status filter on finished records, because that is what most people come here for. Change it to see everything.

Printing a report

Any record prints as a full investigation report — definition, triage, walk-through with the wording that was in force at the time, cause coding, 5 Whys, containment, actions, cost of quality and the review record.

Right First Time — Misblend Investigation Report

Ref: **DONE-1** · WO: **30801** · Misblend 30/08/2026 · Logged 30/08/2026 · Closed 30/08/2026
Generated 14/11/2026, 12:01:37
Coryton Advanced Fuels
DRAFT — INCOMPLETE · 0/35 evidence items reviewed. Outstanding: Investigation walk-through (0/35 evidence items reviewed)

Core Event

Work Order	30801	Serial Number	F26/1801
Product / Blend	Marine Gas Oil DMA	Product Family	—
Failed Parameter	DVPE	Severity	MED
Investigation Owner	S Whitlock	Status	CLOSED
Avoidable	—	Cost of Quality (actual)	£100
Potential Cost of Quality	£200	Rework hours	2 hrs
Issue Summary	Out-of-specification result found during in-process monitoring.		

Investigation Walk-through & Evidence Review

0/35 evidence items reviewed · 0 failures identified

Step 1: Is the result genuine and repeatable? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment
Lab result	—	—	—	—
retest	—	—	—	—
retained sample	—	—	—	—
sample point	—	—	—	—
timing	—	—	—	—

Step 2: Was the correct product, WO, batch and Form B used? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment
Form B	—	—	—	—
WO	—	—	—	—
serial	—	—	—	—
batch allocation	—	—	—	—
pick note	—	—	—	—

Step 3: Do weights, additives, scale checks and transfer/dip records match? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment
Form B sign-offs	—	—	—	—
scale records	—	—	—	—
additive sheets	—	—	—	—
dip book	—	—	—	—

Step 4: Were flushing, draining, deadlegs and stirring completed correctly? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment
Flush record	—	—	—	—
drain-down	—	—	—	—
deadleg risk	—	—	—	—
stirring time	—	—	—	—

Step 5: Could equipment have affected the result? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment
Calibration	—	—	—	—
meter	—	—	—	—
gauge	—	—	—	—
pump	—	—	—	—
temperature	—	—	—	—
valve status	—	—	—	—

Step 6: If production records are correct, was the model/spec/component data correct? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment
Model version	—	—	—	—
component data	—	—	—	—
spec version	—	—	—	—
standing orders	—	—	—	—

Step 7: Was it customer directed, stock unavailable or within test tolerance? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment
Customer instruction	—	—	—	—
stock records	—	—	—	—
tolerance assessment	—	—	—	—

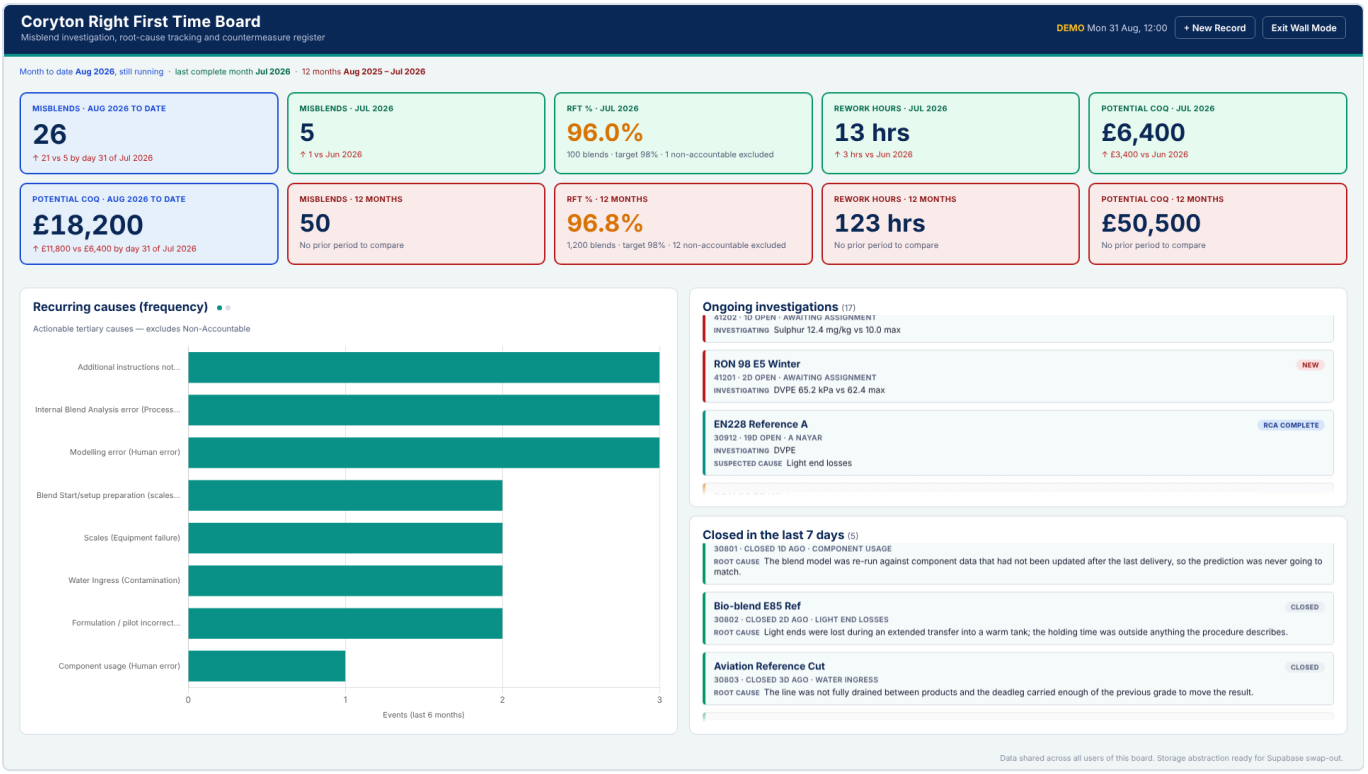
Step 8: Is the same issue recurring? NOT STARTED

Evidence item	Status	Reviewer	Date	Comment

Print to PDF for anything that needs to leave the building.

12. Wall display

Wall Display Mode



The **Wall Display Mode** button turns the board into a passive plant display: no tabs, no modals, cards that do not respond to clicks, bigger type.

It must fit one screen. Nobody walks up to a display and scrolls it, so anything that does not fit rotates instead of stacking. On a narrow screen the ticker joins the rotation and the board shows one panel at a time.

The two lists on the right — ongoing investigations, and closed in the last seven days — scroll continuously. A list that already fits does not scroll, because moving something you can already read is motion for its own sake.

The public wall board

`/wall` is a second page for a display that cannot sign in. It reads an hourly snapshot committed automatically between 08:00 and 18:00 UK time; outside those hours the display stays as it was.

It is the same board — the same tiles, charts and cards, computed by the same rules — and a check renders both pages side by side and fails if any figure disagrees.

That page is served to anyone with the link, with no login, and it carries the ticker detail: blend, work order, owner, causes, root cause, preventive action. That was decided knowingly because the plant display needs it. Costing, 5 Whys and sign-off never appear there.

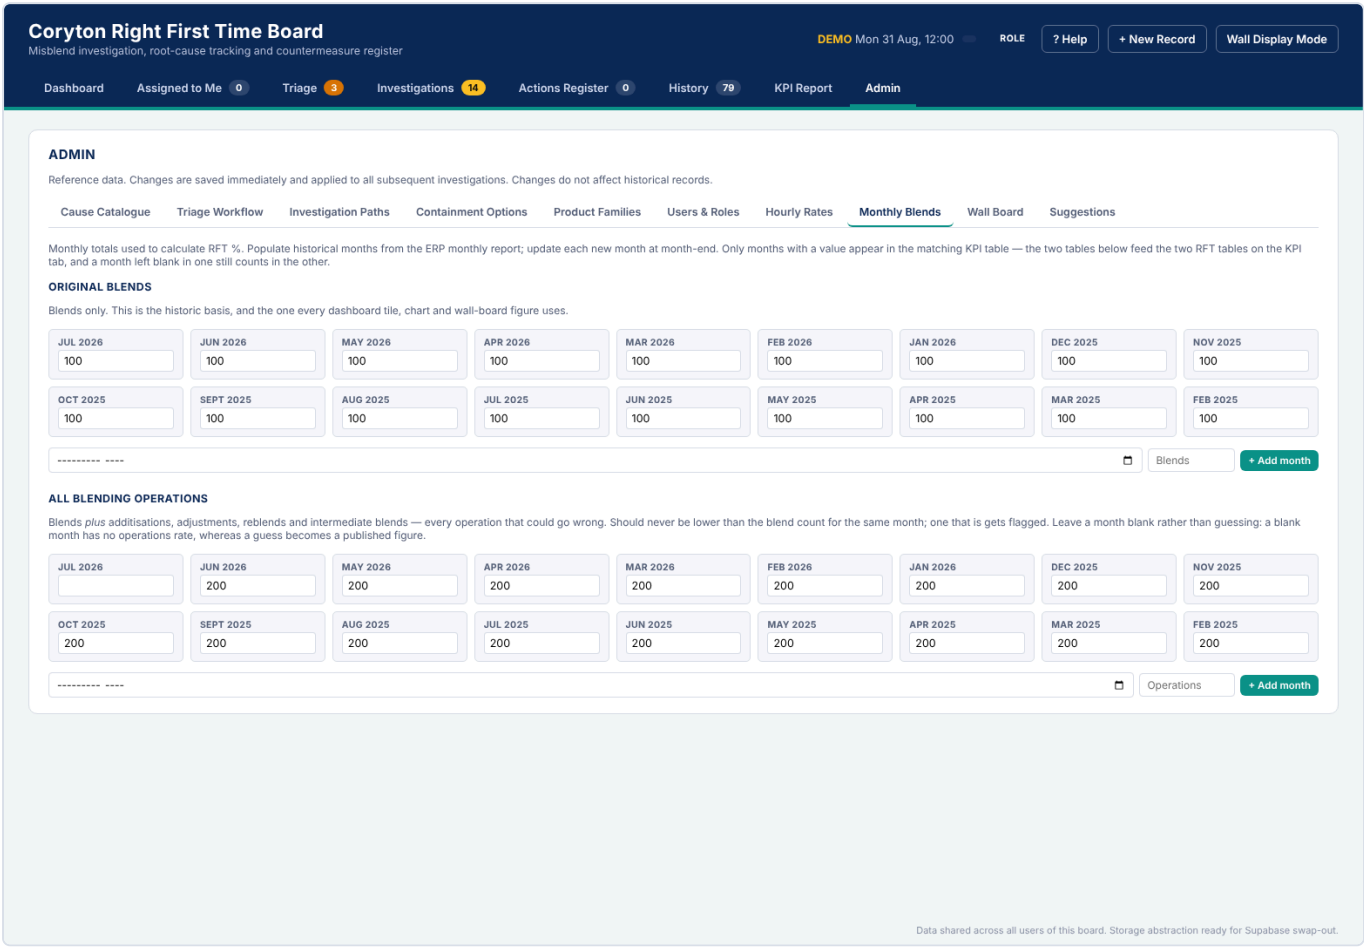
Setting the pace

Admin → Wall Board sets how long a panel stays up and how long a card sits before its turn comes round. Saving restarts the rotation immediately, because whoever is changing it is standing in front of the board judging it.

13. Admin

SECTION	FOR
Cause Catalogue	The taxonomy — codes, areas, and the evidence guide shown against each
Triage Workflow	The T1 checklist wording and its evidence items
Investigation Paths	The fourteen paths, their steps, guidance, evidence and codes
Containment Options	The dropdown on section 5 of a record
Product Families	Used for filtering and cluster analysis
Users & Roles	Everyone with access and their role — to read only; roles are set in the Access app
Hourly Rates	The rates behind the labour half of cost of quality
Monthly Blends	The KPI denominators
Wall Board	Rotation timings
Suggestions	What people have suggested, and what happened to it

Monthly blends — the one entry the KPI cannot do without



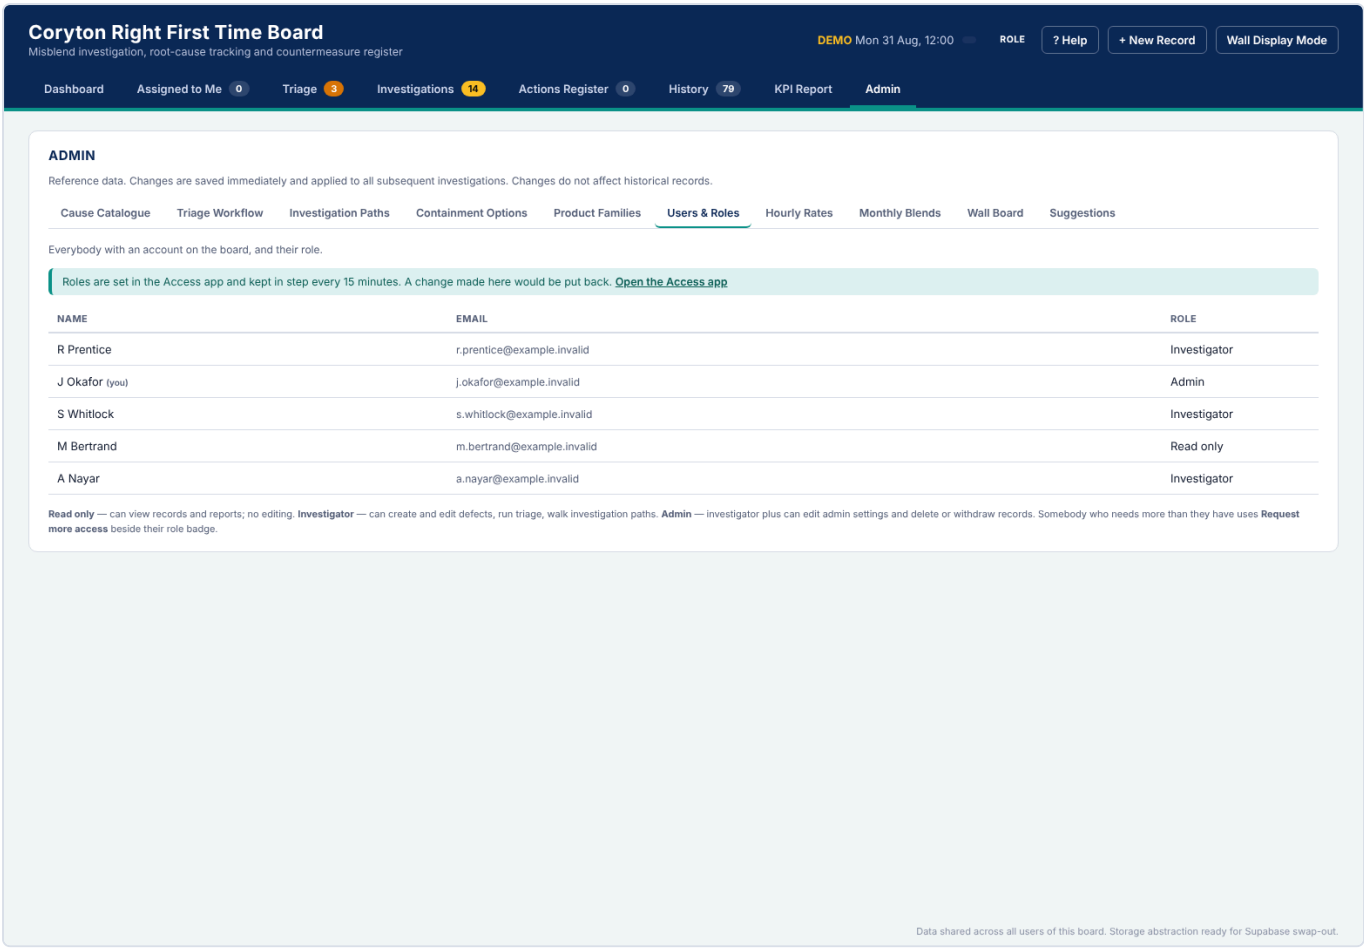
Two figures per month, typed in from the ERP monthly report:

- **Blends** — original blends only. The historic basis, used everywhere.
- **Operations** — blends plus additisations, adjustments, reblends and intermediate blends.

They are typed separately and neither is derived from the other.

Until a month's blend count is entered, that month has no rate anywhere. This is the most common reason somebody reports the KPI looking wrong.

Users and roles



This screen **lists** everybody with an account on the board and their role. It changes nothing. Roles are set in the Access app (<https://coryton-access.netlify.app>) by an Access administrator, and the note at the top says so, with a link that opens Access in a new tab. Access keeps the board in step every 15 minutes, so a role changed anywhere else — in Supabase by hand, say — is put back on the next run. Accounts are made by Access too.

An account with no role set in Access reads as read only. Keep it that way.

Somebody Access has set to *No access* still appears here with the role their account last held: the board cannot see that Access has stopped them signing in. Access is the place to check.

Changes here do not rewrite history

Renaming a cause code, rewording a triage item or rewriting a path step affects **every future investigation only**. Records already walked keep a snapshot of the wording that was in front of the investigator at the time, and print with it.

---

14. Suggesting an improvement

💡 **Suggest an improvement** is in the header on every tab. Say what you were trying to do and what got in the way — the board records which tab you were on, which build you are running and your role, so you do not have to.

Each suggestion gets a reference of the form `RFT-SUG-0001`. The prefix matters: every app in the estate has its own numbering, so the reference says which app it came from.

Suggestions that are taken up become issues in the repository, and the board checks the repository itself to tell you whether yours is genuinely logged — rather than just marking it read.

Several things in this guide exist because somebody raised a suggestion: the misblend date rule, the Misblends column order, and the year-to-date and rolling-12 columns on the KPI table.

---

15. Troubleshooting

Symptom, cause, fix.

A red bar says a save failed

Cause. The board tried to save, the database refused, it retried on a refreshed session, and that failed too. The usual reason is a session that expired while the tab sat idle.

Fix. Press **retry** on the bar. If that keeps failing, check you are still online, then reload and sign in again.

The bar cannot be dismissed, and the record stays open saying it did not save. The board will never tell you a save succeeded when it did not — but **do not close the tab assuming it saved**.

A red bar says someone else saved a newer version

Cause. You changed a record in a tab that had an older copy of it, and somebody else had saved it since. Saving yours would have overwritten theirs, so it was not saved. Nothing of theirs was changed. The bar names the record by work order.

Fix. Note down anything of yours you still need — it is still on screen — then press **Load the latest version**. That shows their version; make your change again on top of it. Other records you saved at the same time were saved normally.

A record should not count

Raised in error, a duplicate, or not a misblend after all: ask an administrator to **withdraw** it. Investigators do not see the button.

An admin opens the record and presses **Withdraw…**, and has to say why. The record stays — in History, under the status filter *WITHDRAWN*, with who withdrew it, when and why at the top — but it is taken out of every count, rate, cost, chart and list, and off both wall boards. It cannot be edited while it is withdrawn. **Reinstate…** puts it back exactly as it was, and the record keeps a note of both.

Delete Entry is for test records only, and is also admin-only: a deleted record is gone for everyone, with nothing to say it existed. Either is only reported once the database confirms it — if it says **NOT withdrawn** or **NOT deleted**, nothing changed.

A yellow bar names a colleague

Cause. They have the record open and hold the edit lock.

Fix. Wait, or ask them. You can read the record and their saves reach you within a few seconds. The lock clears after five minutes without activity.

A banner says a new version is available

Cause. A deploy has landed while you were signed in.

Fix. Apply it at a natural stopping point. Dismissing it is safe — it applies on your next full reload.

A red bar at the top says this is not the live site

Cause. You are on a deploy preview or a local copy.

Fix. Know that **it reads and writes the real database**. There is one Supabase project and every build points at it, so anything you create here is a real record. Use a throwaway record and delete it when you are done.

A month is missing from the KPI report

Cause. Its blend count has not been entered.

Fix. Admin → Monthly Blends. See §10.

A month has a rate on one table and not the other

Cause. It has a blend count but no operations total, or the reverse.

Fix. Enter the missing figure. Neither basis is ever derived from the other, so this is working as intended until somebody types it in.

The RFT % looks higher than the misblend count suggests

Cause. Non-accountable causes are excluded from the rate but not from the count. See appendix A.

Fix. Nothing — but read the note under the tile, which says how many were excluded.

I cannot edit anything

Cause. Your role is read only, which is where every new account starts.

Fix. Click **Request more access** beside your role badge. It opens the Coryton Portal, where you ask for a higher level and an Access administrator decides. Once they have, the change reaches the board within 15 minutes; then **reload the page**. It does not reach you until you do. A board admin cannot change your role from inside the board: roles are set only in the Access app.

Nothing loads

Cause. The board caches its own shell so it opens offline, but every record comes from Supabase.

Fix. Check your connection. If it loads but looks wrong or stale, quote the **Build** badge date from the header when you report it.

Appendix A — what counts against the rate

Cause codes under the area **Non-Accountable** are recorded events that were not a failure of ours — a customer-directed change, a result inside test tolerance, a component the supplier could not deliver. They are still logged, and they still appear in counts and charts. What they do not do is count against the RFT rate.

A record with no cause code counts as ours. An untriaged or uncoded misblend must not quietly improve the figure; that is the wrong way for an unknown to fall.

Counts and rates therefore differ on purpose. The Misblends **tile** shows everything that happened; the **rate** divides by what was ours. Where they differ the tile's note says so — *"· 3 non-accountable excluded"* — so anyone checking the arithmetic finds it right rather than wrong.

Appendix B — 1 August 2026, and the imported records

The board and its investigation methodology **start on 1 August 2026**. Everything dated before that was investigated under the previous process and was imported as archive.

All 569 of those records are **closed**, including the 20 the old register still had marked *Pending Approval*. They were complete investigations awaiting a signature in a process that no longer exists; importing them open would have put them into the sign-off queue to be re-approved under a methodology they were never run under.

They carry no preparer or reviewer name, because nobody signed them off here — being closed reflects **which process they belong to**, not an assertion that anyone approved them. The 20 are marked as having their evidence in the old workbook, so the absent walk-through reads as archive rather than as something somebody forgot to do.

Their status does not affect the figures. Every record counts in the KPI whatever state it is in.