

When the Balance Ties Out but Nobody Can Explain It

How AI-powered OneVue 2.0 turns reconciliation backlogs from a year of work into days


Every finance team knows the uncomfortable version of a reconciliation. The control account ties out at the top. The debits agree with the source receipts. And yet a residual sits there, unexplained, month after month, made up of items nobody has had time to open.

The problem is rarely the arithmetic. It is the volume beneath the arithmetic. A single clearing account can carry a few dozen journal entries sitting on top of several thousand source-system receipts. Somewhere in that population are duplicate postings, split reversals, refunds booked days after the event with the original reference lost, and lump-sum entries that the ERP holds only as a consolidated total with no transaction-level detail behind it. Finding them means reading, sorting, cross-checking and remembering.

So it gets deferred. And deferral compounds. One of the engagements we are working through right now involves clearing close to a year of accumulated transactions. What makes that daunting is not only the hours. It is the confidence problem. When a reviewer works through thousands of lines by hand, the honest answer at the end is often "I think that's all of them." Nobody can prove the population was covered evenly. The judgement was real, but it is not reproducible, and it is not auditable.

Reconciliation is a data problem before it is an intelligence problem

This is where most AI-for-finance conversations start in the wrong place. An assistant that cannot vouch for the completeness of the data underneath it just produces confident-sounding output on an unknown population.



OneVue 2.0 handles that first. Data leaves Oracle Fusion ERP through whichever method suits the volume: a direct REST pull for an on-demand refresh, a scheduled BI Publisher extract for the daily load, Oracle Integration Cloud for event-driven sources, and BI Cloud Connector for establishing multi-year history when a workstream first goes live. All four land in the same place, Ingress Control, and all four are held to the same contract:

- Every load is checked against a control total on record count and value.
- Every load is stamped and versioned as a discrete, auditable snapshot.
- A load that fails or arrives short is routed to IT rather than quietly accepted.

Only then does matching run, and it runs in a deliberate order: deterministic first, then tolerance-based, then judgement. Deterministic rules catch what should always match. Tolerance rules catch what matches once you allow for reality, such as a settlement received net of fees, or a collection made on a Friday that legitimately settles on Monday. Items that clear inside those rules close automatically, with no manual touch at all.

What remains is the hard part. And that residual is exactly where the AI earns its place.


An AI skills engine, not a chatbot

Inside OneVue 2.0, AI is delivered as a library of named, reusable analyses that a reviewer runs against the population behind a specific account and period:

- Duplicate entry pattern watch.
- Settlement timing analysis.
- Refund timing outliers.

- Lump-sum breakdown priority, which works out which underlying records make up a consolidated ERP total.
- GL versus collection residual gap analysis.
- Custom analysis, where a user writes their own prompt and selects the data sources it runs against.

That last point matters more than it first appears. Reconciliation logic is institutional knowledge. It lives in the head of the person who has closed that account for six years. A promptable skill lets that knowledge be written down once, named, and re-run by anyone next month, on next month's data. It stops being a person and becomes a control.





AI Skills Engine
 Configure and execute your analysis.

✕

[< Back to skills](#)

Duplicate Entry Pattern Watch

@ Prompt
 [Re-run](#)

Findings

A few items show a possible duplicate-posting or split-posting pattern, mainly in Bank Deposits and one Courier Interface Journal, where multiple same-day entries share the same batch/reference and appear in close amounts. Most other records look like routine purchase invoice postings or small cash receipts and do not stand out for reversal/credit-close behavior.

Transaction	Finding	Metric	Priority
GL Entry: BANK DEPOSIT Area /1/24-08-2026/223159 Bank Deposits A 223159 11188285 N	This batch appears multiple times on the same date with credits of 30,701.79, 4,913.00, 4,095.30, and 3,908.22, which is more consistent with repeated posting or a split/reversal pattern than a single routine deposit correction.	\$43,618.31 total credits	High
GL Entry: BANK DEPOSIT Area /2/24-08-2026/223159 Bank Deposits A 223159 11188285 N	The same batch/reference also appears twice under Area /2 with credits of 7,557.30 and 694.79 on the same day, suggesting a possible duplicate or follow-on correction rather than an isolated posting.	\$8,252.09 total credits	High
GL Entry: Courier JV /24-08-2026/81267 Courier Interface A 81267 11188299 N	This Courier Interface Journal appears twice on the same date with debits of 621.00 and 663.30, a close same-batch repeat that could indicate a duplicate posting or a reversal/rebook rather than a normal one-off adjustment.	\$42.30 variance	Medium
GL Entry: Payables A 11247629000001 11247642 Y	This payables reference is posted repeatedly across many same-day entries with amounts ranging from 25.00 up to 3,798.00, including several near-duplicate small cash-linked lines; while not a clear reversal, the repetition warrants a quick check for unintended re-entry.	13 postings	Medium

What a finding actually looks like

Run the duplicate-pattern skill on a cash clearing account and you do not get a verdict. You get a ranked list of candidates, each one traceable.

Each finding names the specific journal entry and batch reference, states what the pattern is, and quantifies it. A batch appearing several times on the same date with repeated near-amounts, which is more consistent with a split or reversal pattern than a single routine correction. A courier interface journal posted twice on the same day with a small variance between the two. A payables reference posted repeatedly across many same-day entries, not a clear reversal but enough repetition to warrant a check. Each carries a priority of High or Medium, so the reviewer starts where the money and the risk are.

The language is deliberately careful: possible, suggesting, warrants a quick check. The engine also says plainly which records it looked at and found unremarkable, which is just as useful, because it tells the reviewer where not to spend the afternoon. And every finding keeps the prompt that produced it visible, with a re-run available, so the analysis can be inspected, challenged and repeated rather than trusted on faith.

The line the AI does not cross

READ-ONLY BY DESIGN

OneVue 2.0 has no write access to Oracle Fusion. No API call that posts, no automated journal, nothing. Every path moves data one way only: out of Fusion.

So when a reviewer approves a net-off or a write-off in the workbench, that decision is saved in OneVue as pending. The client's own AP or GL team then enters the adjustment in Fusion exactly as they would any other manual journal. On the next scheduled sync, OneVue pulls that posted journal back in, matches it to the pending decision, and marks it confirmed.

The AI proposes. The reviewer decides. The finance team posts. The system of record stays under the client's control, and the audit trail runs end to end.

From a year of backlog to days of work

Put those pieces together and the shape of the work changes. Deterministic and tolerance matching remove the items that were never really in question. The AI skills engine takes the remaining population and converts it from an undifferentiated list into a ranked, reasoned set of candidates with the evidence attached. The reviewer spends their judgement on decisions rather than on discovery.

In practice, that is the difference between a backlog measured in months and a cleanup measured in days. And, just as importantly, it is the difference between "I think we caught them all" and a documented, repeatable, re-runnable answer.

Reconciliation will never be fully automatic, and it should not be. Someone has to own the write-off. What can change is how much of the work is searching, and how little of it is deciding.

OneVue 2.0's Reconciliation module runs against Oracle Fusion GL, AP and AR, read-only, with configurable workstreams, rules and AI skills. If you are carrying a clearing-account or AP ageing backlog, we would be glad to walk you through it.

