

How we'd defend your numbers

CROSS-INDUSTRY

The method travels. Different operations, same discipline: validate against source, name what breaks silently, hand over the keys. Here's the standard – and the proof it holds across sectors.

The standard

Defensible to the dollar isn't a slogan here – it's a measured bar. On a global SaaS company's revenue rebuild we validated \$84M to 0.002% variance: under \$2,000 of difference across a multi-year migration, every dollar reconciled to source. That is the standard this pack shows you how we hold – on your numbers, in your sector.

0.002%

variance on a rebuilt \$84M revenue pipeline – under \$2,000 total

CRM	\$1,111,111	≠ unreconciled
BILLING	\$1,111,111	
GL	\$1,111,111	

Illustrative reconstruction – anonymized structure, real methodology and totals.

The method, in one page

- 01
- Two independent models, one environment. We build the output two ways – one that matches the legacy system exactly, one clean – and run both in the same environment, so every difference is real, not an artifact of timezones or type casting.
- 02
- Reconcile at every grain. A FULL OUTER JOIN variance check runs at month, then plan, then customer level. A gap that disappears as you drill in is a grain issue; one that persists is a logic error – and now you know which.
- 03
- Parity first, fixes second. We replicate the legacy numbers exactly – bugs included – and validate to near-zero variance before changing anything. Then each bug is fixed in its own reviewed change, where the intent is explicit and signed off.
- 04
- A threshold, not a feeling. The bar was under 0.01%. The result was 0.002% – and the process surfaced 15 silent production bugs that a row-count check would never have caught.

[The full method, including the reconciliation queries →](#)

The questions your board will ask

Can you defend this number to source?

Yes. Every figure traces back through the reconciliation to the system of record. Dual models and a FULL OUTER JOIN at every grain mean we show where each dollar comes from – not just that the total matches.

What breaks silently?

The things a total hides. Parity-first validation is built to surface them – on the \$84M rebuild it found 15 silent production bugs before they reached a board deck, including a \$472K undocumented rate anomaly applied at the database level with no history.

Who controls the logic – pricing, tiers, the rules?

Finance does, in a governed layer it can read – not a hardcoded CASE statement buried in a query nobody owns. When the rules are visible, the anomalies stop being invisible.

What happens when the consultant leaves?

You own it. The work is documented and handed over, and it's senior-led start to finish – no juniors learning on your account.

How fast can we trust the number?

The number is defensible before it's pretty. Parity first means you can sign off on accuracy early; the polish and the fixes come after, in reviewed changes.

Proof for your sector • Cross-Industry

\$84M

validated to 0.002% variance ·
global SaaS revenue

15

silent production bugs surfaced
before the board saw them

Every number in this pack is a real, audited result. Career work – pre-Clarivant, at P&G / eBay – is labelled (career).

\$472K

undocumented rate anomaly, caught
in the audit

100+

locations on one source of truth ·
Carl's Jr

Bring us the number nobody trusts.

A 30-minute call, direct with the founder – no pitch. We'll tell you, on your own data, exactly how we'd defend it.

Book a 30-min strategy call

clarivant.ai