

DATA ENGINEERING

Your dbt tests can pass every CI run — and still miss the exact failure that reaches production.



THE PROBLEM

**A freshness or time-window test proves one thing:
the data looked right at the moment CI ran.**

It says nothing about whether it stayed right.



THE REAL PATTERN

Time-specific tests pass cleanly in CI/CD.

Then fail in production — because the data is time-sensitive, and nobody's re-checking it between test runs.



WHY IT SLIPS THROUGH

CI runs on its own schedule.

Your pipeline breaks on its own schedule.

Those two schedules don't have to line up —
and usually don't.



THE CONCRETE FAILURE MODE

A 24-hour freshness check passes at 9am.

The pipeline stalls at 10am.

Nothing re-checks until tomorrow's run.

**That's 23 hours of silent staleness —
fully “tested.”**



WHY AI MAKES THIS WORSE, NOT BETTER

AI can generate a full dbt test suite in minutes.

Structurally valid. Technically correct.

Easy to ship.

It writes to the same blind spot: tests that
prove existence, not continuous coverage.



WHAT DBT LABS' 2026 REPORT CONFIRMS

“AI-driven acceleration is outpacing trust and governance.”

— dbt Labs, State of Analytics Engineering 2026

More tests, generated faster, isn't the same as more protection.



THE REAL FIX

Before shipping a test, ask: does this prove things looked fine when I checked — or does it run often enough to catch the failure that actually matters?

Test cadence is the missing dimension, not test count.

Your turn

What's the shortest gap you've seen between “test passed” and “pipeline broke” in production?

Follow for more on data quality, dbt, and where AI actually helps in the pipeline.