

Every Row Has a Destination

Reconcile one-to-one data moves before calling a migration complete

ENGLISH · 1.0 · OCTOBER 2026

SAMPLE

Unpublished review edition. New product terms and checkout are not approved.

Start with one bounded decision

An import screen says 96 rows succeeded. That number is hard to interpret without the expected source set and the records deliberately rejected. This kit starts from a frozen source list and asks a narrower question: did every source row receive one documented disposition? It will not prove that every field was mapped correctly. Its value is making omissions and contradictory dispositions easy to inspect.

People reviewing a CSV migration in which each source row must receive exactly one final disposition.

Use a row ledger, counts and a local read-only checker to find duplicate dispositions, missing source rows and unexpected destination rows before a one-to-one migration handoff.

One method to try

Freeze a source identity

Give each source row a stable migration ID before splitting files or correcting values. A business key can repeat; a row ID must identify this row in this frozen run. Keep a source fingerprint in the decision sheet. If source rows change, use a new run ID. Otherwise a clean report can refer to different inputs than the team actually imported.

Sample boundary

This sample stops here. The full review package contains the input dictionary, worked case, calculator instructions, controlled-change exercise, answer key and reusable CSV templates.