

CHECKLIST

Data Migration Readiness Checklist

Six stages and 31 checks that decide whether a migration lands cleanly or gets fixed in production.

How to use this

Migrations rarely fail in the build. They fail because a stage was skipped or thinned under schedule pressure, and the gap shows up at cutover or after go-live, when it costs the most to fix. Work through this list before you commit to a go-live date, and again at the end of each stage. Any item you cannot answer with a document, a date or a number is where your risk is.

Each stage ends with a gate: the evidence is reviewed, and someone with the authority to stop the project decides go or no-go. A gate that cannot stop the project is decoration.

AUTHOR

Dat Tran, Enterprise Data & AI Analytics Architect

EDITION

October 2026

01 Scope and profiling

Why it matters: estimates built on a count of tables rather than on profiled data are where migration schedules usually break, at cleansing or user acceptance testing, after the date is already committed.

- ☐ Scope agreed in writing: which tables are migrated, which are archived, which are dropped.
- ☐ Source data profiled at full volume (volumes, nulls, patterns, orphan records) before the estimate is signed off.
- ☐ Every figure in the plan traceable to the profiling report.
- ☐ Project plan covers every stage through to decommissioning, each with an owner and a date.
- ☐ Measurable exit criteria defined for each stage before that stage starts.
- ☐ RAID log (risks, assumptions, issues, dependencies) opened in week one and reviewed weekly.

02 Mapping

Why it matters: transformation rules written from documentation or assumption surface as misclassified records after go-live, when the fix means correcting production data and rebuilding users' trust.

- ☐ Source-to-target mapping written at field level, not table level: type conversions, code mappings, conditional logic, null handling.
- ☐ Every rule checked against the values actually found in profiling, not the values listed in the data dictionary.
- ☐ Rules for derived values signed off by a business subject-matter expert, not decided by the engineer.
- ☐ Unmapped fields in both directions brought to zero by explicit decision.
- ☐ Business sign-off on the mapping specification, with no rows left as "TBD".

03 Data quality

Why it matters: the source keeps changing after you profile it. Without a re-check, the data you tested is not the data you cut over.

- ☐ Quality thresholds for the source data agreed with the business before cleansing starts.
- ☐ Data quality metrics meet the agreed thresholds, and the issue log is closed.
- ☐ Source re-profiled before cutover to catch new values and new volume.
- ☐ Delta-load plan in place for changes made to the source between the main load and cutover.

04 Test and reconcile

Why it matters: a job that finished without errors is not evidence that the data is right. Tests on samples miss rare historical values and the effects of scale, and without automated reconciliation, differences are found by users weeks later with nothing left to trace them back to.

- ☐ Migration pipeline runs end to end in a test environment, with reconciliation built in.
- ☐ Load tested at full volume on representative infrastructure, including failure scenarios such as kill-and-restart and a duplicate run.
- ☐ Automated reconciliation runs after every load, against variance thresholds agreed in advance, with results kept for each round.
- ☐ A full reconciliation re-runs with one command, not a week of spreadsheet work.
- ☐ Unit, system integration, user acceptance, volume and reconciliation tests passed, with no open critical defects.

05 Cutover

Why it matters: without numeric triggers agreed in advance, the choice between pushing on and rolling back is made under pressure in the middle of the night, and sunk cost usually wins.

- ☐ Cutover window agreed with the business.
- ☐ At least two full-volume mock cutovers run, with timings measured and the runbook verified by the people who will execute it.
- ☐ Rollback triggers set as numbers (a hard-stop time and a variance threshold) and signed by the sponsor before go-live.
- ☐ Source protected from writes throughout the load, by design rather than by instruction.
- ☐ Rollback path rehearsed at least once.
- ☐ Go-live decided against criteria agreed beforehand: reconciliation passed and business sign-off given.

06 After go-live

Why it matters: a migration ends when the old system is switched off, not at go-live. Until then you pay for two systems, and some users keep entering data in the old one.

- ☐ Hypercare team named, with exit criteria signed off: incident rate back to normal, reconciliation stable.
- ☐ Hypercare ends when its exit criteria are met, not when the calendar says so.
- ☐ Every temporary coexistence mechanism, such as an interim sync, has an end date.
- ☐ Decommission date set in the business case and tracked as a milestone.
- ☐ Legacy system switched off only when no consumer still reads from it and the archive has been verified against retention rules.

Five questions for any migration plan

When I have half an hour to review a migration plan, these are the questions I ask. A sound plan answers with numbers and dates. A risky one answers with "we will".

- 1 Which profiling report is this estimate based on?
- 2 Does the rule for each coded field cover every value in the data, or only the values in the documentation?
- 3 Which full-volume run do the runbook timings come from?
- 4 At what variance, and by what time, does someone call a rollback, and who is it?
- 5 How long does a full reconciliation re-run take: one command, or a week?

Integration & Migration Review

If a migration is coming and you want an independent read before you commit to a date, the Integration & Migration Review is a fixed-scope review built for that point. EthanCorp designs and runs data integration, analytics and AI automation systems for organisations that need their fragmented data to be trustworthy, built as durable systems, not demos. The review profiles your source systems as they actually run today, ranks the failure points by business impact, and produces a first source-to-target mapping for the domain that matters most.

[Book a review →](#)