

SQL Server → PostgreSQL Migration Pack

Five phases, sample sequencing — calibrate durations with your pilot, not this document.

Pack	SQL Server → PostgreSQL Migration Pack
Document	Sample project plan (5 phases)
Date	2026-09

AnovaCloud — Engineering the AI-Native Enterprise. Frisco, Texas. This pack is a working instrument, not a brochure: complete it with your team, score honestly, and bring it to your discovery call. © 2026 AnovaCloud.

HOW TO USE

How to use

Durations are illustrative samples. Run the pilot migration first; pilot actuals calibrate the program estimate. Owners are roles — assign names before kickoff.

PHASE 1 — ASSESS (SAMPLE: WEEKS 1-3)

Phase 1 — Assess (sample: weeks 1-3)

Owner: DBA lead + Architect

1. Estate inventory & feature scan

Complete instance/database inventory; automated feature scan per database. Deliverable: Inventory workbook + scan output

2. Application dependency mapping

Dependency questionnaire with every app owner. Deliverable: Signed dependency map

3. Compatibility scoring

Per-object T-SQL/CLR/SSIS scoring; complexity per database. Deliverable: Compatibility report

4. Business case & TCO

3-year stay-vs-migrate model reviewed with finance. Deliverable: Signed business case

5. Wave plan v1

Sequenced waves: easy wins first; entry/exit criteria per wave. Deliverable: Wave plan

PHASE 2 — CONVERT (SAMPLE: WEEKS 4-8)

Phase 2 — Convert (sample: weeks 4-8)

Owner: Migration engineers

1. Schema conversion pilot

Convert one representative schema end to end with tooling. Deliverable: Converted schema + tool report

2. T-SQL remediation backlog

Hand-conversion items logged with owner and estimate. Deliverable: Backlog

3. Data-type mapping sign-off

Mappings reviewed with application owners. Deliverable: Signed mapping

4. App remediation

Application SQL/ORM changes per the dependency map. Deliverable: Merged app changes

5. SSIS/Agent disposition

Packages and jobs rewritten or re-platformed. Deliverable: Working replacements

PHASE 3 — MIGRATE (SAMPLE: WEEKS 6-12, OVERLAPS)

Phase 3 — Migrate (sample: weeks 6-12, overlaps)

Owner: Data engineers

1. Load + CDC per wave

Full load then continuous replication; lag SLO monitored. Deliverable: In-sync target

2. Validation harness

Row counts, checksums, sampling automated per wave. Deliverable: Green validation report

3. Performance replay

Top queries replayed vs. baseline; regressions owned. Deliverable: Performance sign-off

4. Security & HA build-out

Roles, RLS, auditing, HA topology, backups live on target. Deliverable: Operational target

5. Cutover rehearsal

Full rehearsal in non-prod with timings. Deliverable: Rehearsal report

PHASE 4 — CUTOVER (PER WAVE)

Phase 4 — Cutover (per wave)

Owner: Cutover commander

1. Go/no-go gate

Objective criteria: lag, validation, smoke tests. Deliverable: Go decision

2. Connection flip

Sequenced DNS/config flip; smoke tests green. Deliverable: Live on PostgreSQL

3. Hypercare

48-hour watch with escalation. Deliverable: Stable operations

4. Rollback readiness

Rollback rehearsed; triggers and authority named. Deliverable: Rollback plan

PHASE 5 — OPERATE

Phase 5 — Operate

Owner: Platform team

1. Runbooks & monitoring

PostgreSQL-native runbooks, dashboards, alerts. Deliverable: Ops handover

2. Decommission source

Per-wave source retirement; license true-down. Deliverable: Decommissioned estate

3. Team enablement

Training and on-call shadowing complete. Deliverable: Self-sufficient team

4. Program retrospective

Actuals vs. estimate; lessons recorded. Deliverable: Retrospective