Faysal Aziz

Break Free: Enterprise Database Migration

Strategies for Migrating Oracle and SQL Server to PostgreSQL - 1 day

FA-0192 · 1 day

Data & Analytics

Introduction

Migrating mission-critical databases from proprietary engines like Oracle and Microsoft SQL Server to PostgreSQL has evolved from a cost-cutting tactic into a strategic imperative for enterprise agility. However, successful schema transformation and data movement require navigating significant architectural differences in concurrency models, procedural logic, and data types.

This intensive one-day course equips engineering teams and database leaders with proven techniques to plan, execute, and validate database migrations while minimizing downtime and business risk.

Led by an instructor with over 30 years of industry experience, the course prioritizes battle-tested, production-grade migration patterns over theoretical frameworks. Participants will explore practical schema conversion using modern automation toolchains (such as AWS SCT, ora2pg, and pgloader), PL/SQL and T-SQL refactoring, change data capture (CDC), automated validation testing, and post-migration performance optimization.

Learning outcomes

- Assess source databases and map complex proprietary features to native PostgreSQL equivalents
- Convert relational schemas, data types, and constraints accurately while preserving data integrity
- Refactor PL/SQL and T-SQL stored procedures, functions, and triggers into PL/pgSQL
- Utilize automated conversion toolchains alongside manual code refactoring strategies
- Design low-downtime and zero-downtime migration strategies using Change Data Capture (CDC) and logical replication
- Plan and execute physical data extraction, transformation, and bulk loading workflows

- Validate data consistency and business logic fidelity across source and target environments
- Identify and resolve query execution discrepancies arising from engine-level optimizer differences
- Establish post-migration operational, monitoring, and fallback procedures

Prerequisites

- Hands-on experience administering or developing on Oracle, SQL Server, or PostgreSQL databases
- Solid understanding of relational database concepts, SQL queries, indexing, and transactions
- Basic familiarity with command-line operations, shell environments, and database configuration files

Target Audience

- Database Administrators and Database Engineers executing cross-engine migrations
- Enterprise Architects and Solutions Architects designing modern data platform strategies
- Software Engineers and Backend Developers responsible for query and stored procedure refactoring
- Infrastructure, Platform, and DevOps Leads planning database modernization initiatives

Training outline

1. Migration Assessment and Architectural Mapping

Pre-migration assessment and discovery

1. Source workload cataloging and schema dependency mapping
2. Feature compatibility matrix and technical gap analysis
3. Estimating migration effort, complexity, and risk factors
4. Database sizing, licensing, and TCO evaluations

Architectural divergence: Proprietary vs. PostgreSQL

1. Concurrency control and transaction isolation model differences
2. Process model vs. thread-based execution models
3. Storage architecture, block sizing, and tablespace concepts
4. Proprietary features vs. open-source alternatives

2. Schema and Data Type Conversion

Schema conversion toolchains and strategies

1. Automated tools: AWS SCT, ora2pg, pgloader, and native utilities
2. Manual refactoring boundaries and rule-based transformations
3. Object mapping: tables, views, sequences, and synonyms
4. Primary keys, foreign keys, and referential constraints
5. Index translation: B-Tree, partial, functional, and specialized indexes

Data type translation and handling

1. Numeric and character type mappings
2. Date, time, and timezone representation differences
3. LOBs, CLOBs, BLOBs, and VARBINARY conversion patterns

3. Code Conversion and Business Logic Refactoring

Stored procedure and function translation

1. PL/SQL and T-SQL to PL/pgSQL language translation
2. Exception handling mechanisms and error code mappings
3. Package decomposition into PostgreSQL schemas and modules
4. Autonomous transactions and dynamic SQL refactoring

Triggers and dialect adaptation

1. Row-level and statement-level trigger differences
2. Proprietary SQL syntax adaptation to standard ANSI SQL
3. Hierarchical query conversion using Recursive CTEs
4. Proprietary built-in functions and custom helper function creation

4. Data Migration and Synchronization Strategies

Migration execution patterns

1. Offline batch migration vs. online low-downtime migration

2. Big-bang cutover vs. phased or incremental cutover
3. Fallback, rollback, and dual-write architectural patterns
4. Target environment provisioning and baseline parameter tuning

Data extraction, movement, and loading

1. High-throughput bulk export and import workflows
2. Parallel data extraction and parallel COPY loading

5. Validation, Tuning, and Post-Cutover Operations

Data verification and regression testing

1. Row count, checksum, and data integrity verification
2. Automated query result and business logic validation
3. Schema structure and constraint verification routines
4. Performance benchmarking and baseline comparison

Cutover execution and operational readiness

1. Cutover runbook formulation and execution timeline
2. Production switchover, DNS updates, and application re-pointing
3. Post-migration monitoring, lock contention, and health checks
4. Backup strategy implementation and operational handoff

Disclaimer

This course outline is provided as a general training framework and proposed scope of coverage only. It does not constitute a fixed agenda, binding commitment, official certification pathway, partnership representation, or guaranteed delivery sequence. The trainer reserves the right, at his professional discretion, to amend, expand, condense, reorder, replace, or otherwise modify any portion of the content, emphasis, structure, tools, demonstrations, or delivery approach without prior notice, based on participant background, organizational requirements, technology changes, platform availability, security considerations, and instructional judgment.

© 2026 Faysal Aziz · Kuala Lumpur