

Db2 to PostgreSQL: Zero Downtime Database Modernization

Migrating a global HR leader's complex on-premises IBM Db2 database to a fully managed, cost-effective AWS RDS for PostgreSQL service, ensuring zero downtime.

Overview

A global HR services leader sought to align with their cloud strategy by migrating their on-premises IBM Db2 database (approx. 500 GB across 600 tables) to AWS RDS for PostgreSQL (v16.2).

- Delivered a seamless, near zero-downtime migration from the end-of-life Db2 version to a fully managed, cost-effective PostgreSQL cloud service.
- Ensured full schema compatibility and data integrity by successfully navigating complex encoding mismatches (IBM-943 to UTF-8) and performing custom data validation.
- Reduced operational risk and high licensing costs by moving to an open-source, cloud-native solution on AWS, setting the stage for enhanced scalability and efficiency.



Client Profile

Based in Japan, our client is a global leader in human resource services providing flexible HR solutions tailored to the evolving labor market.

Challenges: Mitigating Licensing Costs and EOL Risks

The client faced significant financial and operational hurdles that required a swift resolution

- **Prohibitive IBM Db2 Licensing Costs:** The high cost of proprietary database licenses was a major inhibitor to the client's cost-efficiency goals.
- **End-of-Life (EOL) Threat:** IBM was planning to discontinue support for Db2 versions up to 10.5 within a year, posing a critical operational risk to the existing database functionality (v9.7).

- **Complex Data & Schema:** Migrating approximately 500GB of data and 600 tables, including complex Db2-specific triggers and functions, while maintaining zero data loss.
- **Encoding Mismatch:** Overcoming initial challenges with the original IBM-943 encoding and its compatibility with the target UTF-8 encoding in PostgreSQL.

Solution: Cloud-Native Migration to PostgreSQL

QBurst proposed and executed a robust migration strategy to move the on-premises IBM Db2 (v9.7) to AWS RDS for PostgreSQL (v16.2). This solution leveraged PostgreSQL's strong features and open-source cost-effectiveness while utilizing AWS cloud services for a fully managed experience.

The migration involved a multi-step approach:

- The legacy Db2 schema was converted to PostgreSQL using the AWS Schema Conversion Tool (SCT).
- Data was extracted from Db2 (v9.7), temporarily imported to an intermediate Db2 11.5 instance (created with IBM-943 encoding to resolve compatibility issues), and then moved.
- The final data transfer to AWS RDS PostgreSQL was handled using AWS Database Migration Service (DMS), utilizing shell scripts and utility commands to process the large 500GB dataset in optimized batches.
- Manual adjustments were essential for complex Db2 triggers, sequences, and functions that did not convert directly, ensuring the application logic remained intact in the PostgreSQL environment.

Technical Highlights

To ensure speed, integrity, and compatibility throughout the migration process:

- Migrated from Db2 v9.7 to AWS RDS PostgreSQL v16.2.
- Used Db2 Load commands over Import commands to optimize large batch processing and reduce log-related crashes.

- Employed IBM Data Studio and AWS SCT for DDL extraction and schema conversion.
- Implemented a custom, character-length based data validation strategy to bypass AWS DMS failures caused by encoding differences (Chinese characters, IBM-943 vs. UTF-8), ensuring data integrity.
- Disabled PostgreSQL triggers during DMS migration to prevent data inconsistencies.

Impact: Ensuring a Smooth and Hassle Free Migration

The successful migration delivered immediate and long-term strategic benefits:

- **Significant Cost Reduction:** Elimination of high IBM Db2 proprietary licensing and maintenance fees.
- **Zero Downtime:** Achieved continuous database functionality throughout the migration process.
- **Future-Proofing Data Infrastructure:** Shifted to a fully managed, EOL-resistant, cloud-native AWS RDS PostgreSQL solution.
- **Enhanced Scalability and Performance:** Optimized data import scripts using batch processing, ensuring faster, reliable migration of the 500GB dataset.
- **Mitigated Operational Risk:** Addressed the threat of IBM EOL support, securing uninterrupted service functionality.