

Multiple heterogeneous databases migrations

AWS Data Migration Services Case Study

Executive Summary

Brightfield, with its Augmented Analytics Platform named Talent Data Exchange (TDX), provides a suite of analytical applications that help optimize both, extended workforce, and employee talent segments.

Brightfield seeks to begin their migration to the cloud, taking advantage of the benefits of a platform like AWS. Also, the approach considers the practice of infrastructure as code (IaC) to automate deployments. Brightfield had a significant investment in Oracle databases, and rather than refresh the hardware and renew the licenses, they opted to migrate the databases to Amazon Web Services.

The Challenge

The client has multiple large on-premises Oracle databases on legacy hardware, using a data center provider to host them. The client's IT strategy was to find a way to migrate their information to Amazon Web Services that would allow them to have full control over their data. The database size was large (over 15TB), and the workload was highly transactional, both of which complicated migration planning. A more complex migration strategy was required to accommodate the nature of the platform.

Why AWS

AWS Database Migration Service (DMS) provides a fully managed cloud service for the more natural migration of relational databases, NoSQL databases, data warehouses, and various types of data stores.

AWS DMS can help you with the migration of your data to the AWS Cloud, between the combinations of on-premises and cloud setups, or between on-premises instances by leveraging an AWS Cloud setup. Users can easily replicate ongoing modifications for maintaining synchronization of sources and targets as well as perform one-time integrations efficiently. The AWS Schema Conversion Tool (SCT) is ideal for scenarios where you want to switch database engines before migrating data using the AWS Database Migration Service tool.

About Customer



Brightfield is an Artificial Intelligence and Big Data Analytics company that optimizes contract labor and performance for employers and staffing firms.

“Database freedom on AWS.”

Migrating to AWS has cut Amazon's annual database operating costs by more than half, despite having provisioned higher capacity after the move. Database-administration and hardware-management overhead have been significantly reduced, and cost allocation across teams is much simpler than before.

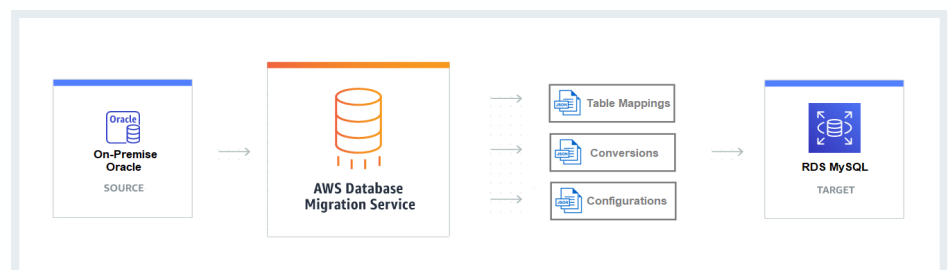
Data migration using AWS DMS also ensures the benefits of higher flexibility, security, cost-effectiveness, and faster speed to market, which come with all AWS services.

The Solution

Multiple Proof of Concept (POC) approaches were evaluated to determine the best solution to transfer the data between the data center and AWS. Considering the client's needs and the improvements that are to be made to the databases, which include converting the name of the databases and their fields to lowercase and the restriction of schedules to be able to do the data migration, it was necessary to use of all the faculties that AWS offers:

- AWS MySQL RDS as a target database to simplify the overall management and data replication of the database.
- Utilization of AWS Data Migration Service (DMS) to migrate the on-premises database with an Oracle instance.
- Implementation of custom scripts in JSON to change the name of the databases and their fields to lowercase and additional data validation.
- Create a robust architecture that can support and give better agility to data migration by AWS DMS (Multiple instances replication with the maximum levels of computation, network, memory capacity, maximum number of tables to load in parallel, maximum commit rate during full load, create multiple tasks of migration divided depending on the size of the tables of the DB, identification of tables that contain LOB type fields to configure suitable jobs for their replication).

After standardizing on the legacy Oracle schema of 376 main tables (PK's, FK's, Views, Functions, and Index), the team adopted a three-step migration approach. First, the team built and ran functional tests on an Amazon RDS-based prototype for MySQL to ensure that all conversions performed by the AWS Schema Conversion Tool (AWS SCT) and AWS Database Migration Service (AWS DMS) would be successful. Second, on a pre-production database, the team ran application API latency tests and fine-tuned the AWS DMS settings based on production data. In the third and final step, the team created and tested a rollback plan and data validation technique that would allow recovery in the event of any problems.



Best features of AWS RDS

- Easy to administer
- Highly scalable
- Inexpensive
- Secure
- Available and durable

Results and Benefits

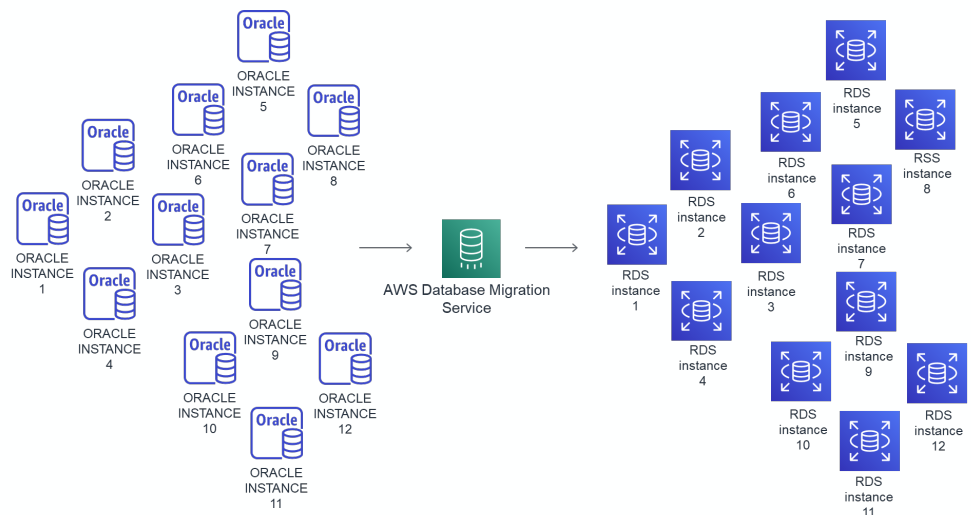
By switching from a local Oracle database solution to Amazon RDS that supports MySQL, Brightfield gained considerable capabilities for scalability, performance, and fault tolerance. Oracle's change significantly reduced checkpoints and the resulting network consumption. By moving to Amazon RDS, performance doubled.

Because Amazon RDS for MySQL offers such extensive automation, maintenance is now much easier and faster.

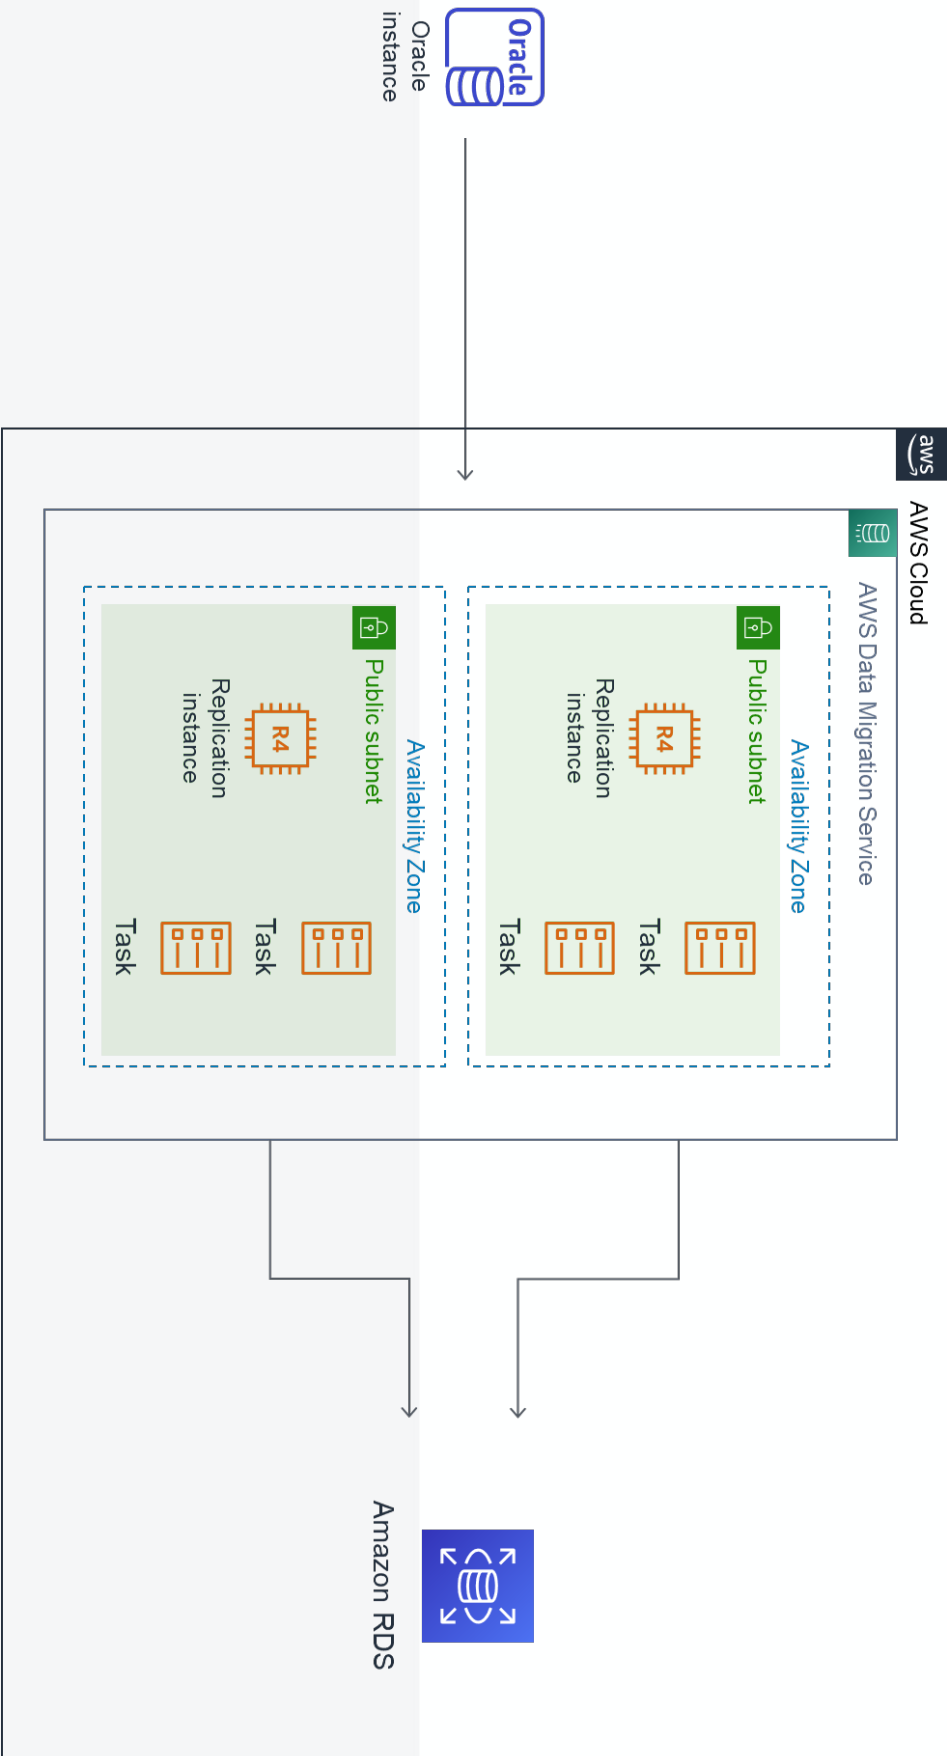
Brightfield is enjoying these benefits with cost and time savings. Also, new monitoring systems for the databases were implemented easily, thus avoiding the time it would have taken to build them from scratch. Amazon RDS for MySQL automatically adds storage as needed, which is helpful in the case of unexpected spikes, since you only pay for what is used.

The migration was completed with the following results:

- The cutting activities were executed in just minutes since DMS fully migrated the data in no more than 5 hours per DB, executing up to 3 DB in parallel.
- AWS Data Migration Services allowed database migration to take place "behind the scenes" with less customer resources and without interfering in daily operations and workflows.



DMS Solution Diagram



Next Steps

AWS provides building blocks that you can assemble quickly to support virtually any workload. With AWS, you will find a complete set of highly available services that are designed to work together to build sophisticated, scalable applications.

You have access to highly durable storage, low-cost computing, high-performance databases, management tools, and more. All this is available without up-front costs, and you only pay for what you use. These services help organizations move faster, lower IT costs, and scale. AWS is trusted by the largest enterprises and the hottest start-ups to power a wide variety of workloads, including web and mobile applications, game development, data processing and warehousing, storage, archive, and many others.

Superior Performance

This infrastructure provides a fast, resilient, and high availability environment for the application.

LOW TCO

Save money by replacing physical hardware with expensive license fees, with AWS you pay for what you use.

Fully Managed

With fully managed resource provisioning, maintenance, and backup, deployments are more efficient.

Minimal downtime

AWS Database Migration Service helps you migrate your databases to AWS with virtually no downtime. All data changes to the source database that occur during the migration are continuously replicated to the target, allowing the source database to be fully operational during the migration process.

Reliable

The AWS Database Migration Service is highly resilient and self-healing. It continuously monitors source and target databases, network connectivity, and the replication instance.

About IO Connect Services

IO Connect Services is a company specializing in Information Technology Consultancy Services. All our team members have one thing in common: our enthusiasm for technology and our passion for customer service excellence. We provide services in all North America, LATAM and Europe. Our headquarters are in NYC metropolitan area, and we also have offices in Guadalajara, Mexico and Madrid, Spain.