



Mindtree

A Larsen & Toubro Group Company



# OSS to Azure Cloud<sup>®</sup> Database Migration: A Mindtree Approach



# CONTENTS

|   |   |
|---|---|
| ↔ Introduction .....  | 3 |
| ↔ High Level Architecture .....   | 3 |
| ↔ Migration Strategy .....  | 4 |
| ↔ Migration Framework .....   | 5 |
| ↔ OSS Use Cases (RDBMS to Azure) .....                                  | 6 |
| ↔ Case Study: Oracle to Azure PostgreSQL Using Compass and Ora2pg ..... | 7 |
| ↔ Recommended Tools for Migration (Compass) .....                       | 8 |
| ↔ Recommended Tools for Migration (Ora2pg) .....                        | 9 |
| ↔ Benefits .....  | 9 |

# Introduction

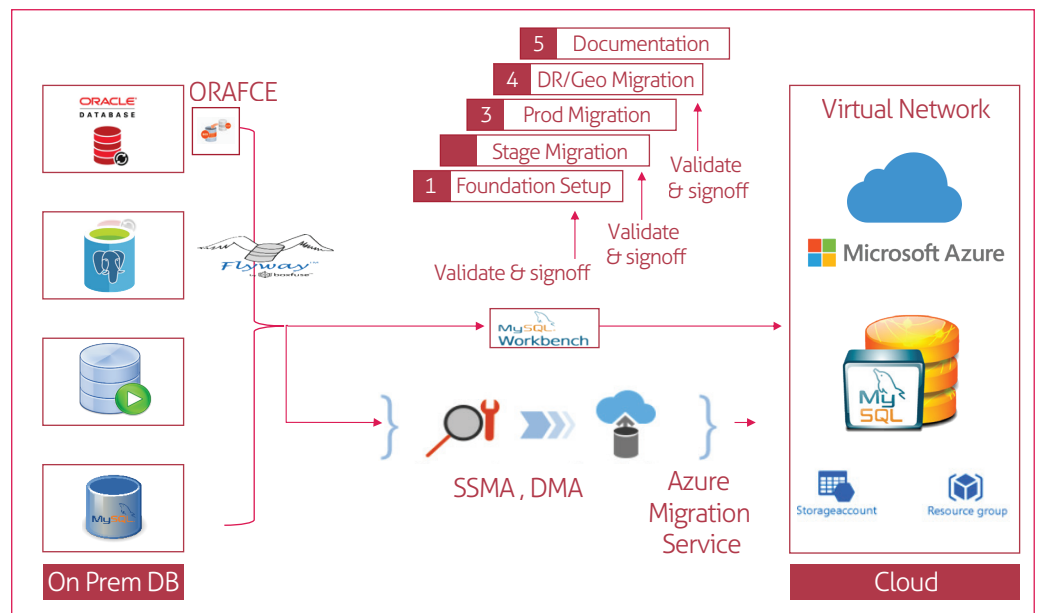
Gone are the days of large-scale database systems that offer more than a thousand functionalities. Customers use less than 20% of these functionalities provided by the specific DBMS/RDBMS systems, and hence, end users are moving towards database systems which offer focused services. It has also been observed that in case an end customer is to integrate more than two-three small-scale database systems for multiple functionalities, it becomes easier and cost effective. Interestingly, most such databases are open source in nature. At the same time, the trend to move things to cloud has also been established as a standard.

In this whitepaper, we have attempted at summarizing a few ways of migrating some of the major open source databases and enterprise databases (like Oracle) to cloud database (predominantly PaaS).

## High Level Architecture

The overall high-level architecture follows the process of Assess, Migration, and Optimize to ensure that workloads are ready to meet production demands on the Azure platform.

A high level solution architect is flexible enough to cater to various different RDBMS (heterogeneous) needs by using tools like ORA2PG .DMA SSMA.



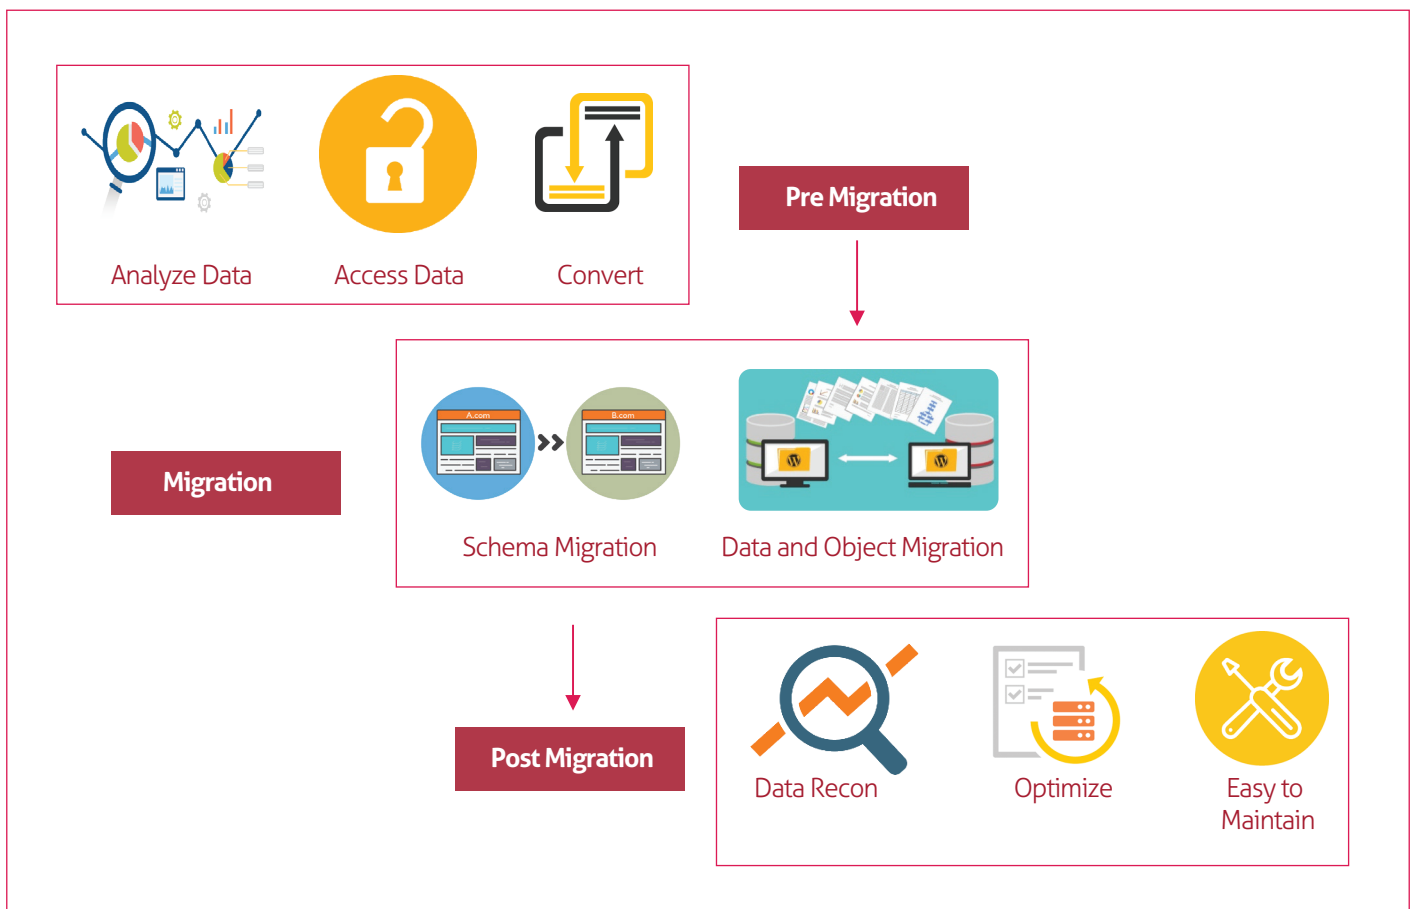
Our approach to database migration, in general, could be summarized as:

- Pre-migration check-list preparation (Install DMA, assess database compatibility)
- Offline versus online migrations to determine whether the downtime is acceptable; if not, to be prepared for online/minimal down time migration options (DMS, Transactional Replication)
- Preparation of target Azure Postgre Database
- Validate the migration steps
  - Create a migration project
  - Analyze and access source and target
  - Add databases and plan
  - Select login
  - Migration to Azure Postgre Database
  - Verify data in Azure Postgre Database
- Post-migration: Optimum utilization and to ensure smooth functioning wherever possible
- Remediate procedures/ steps

# Migration Strategy

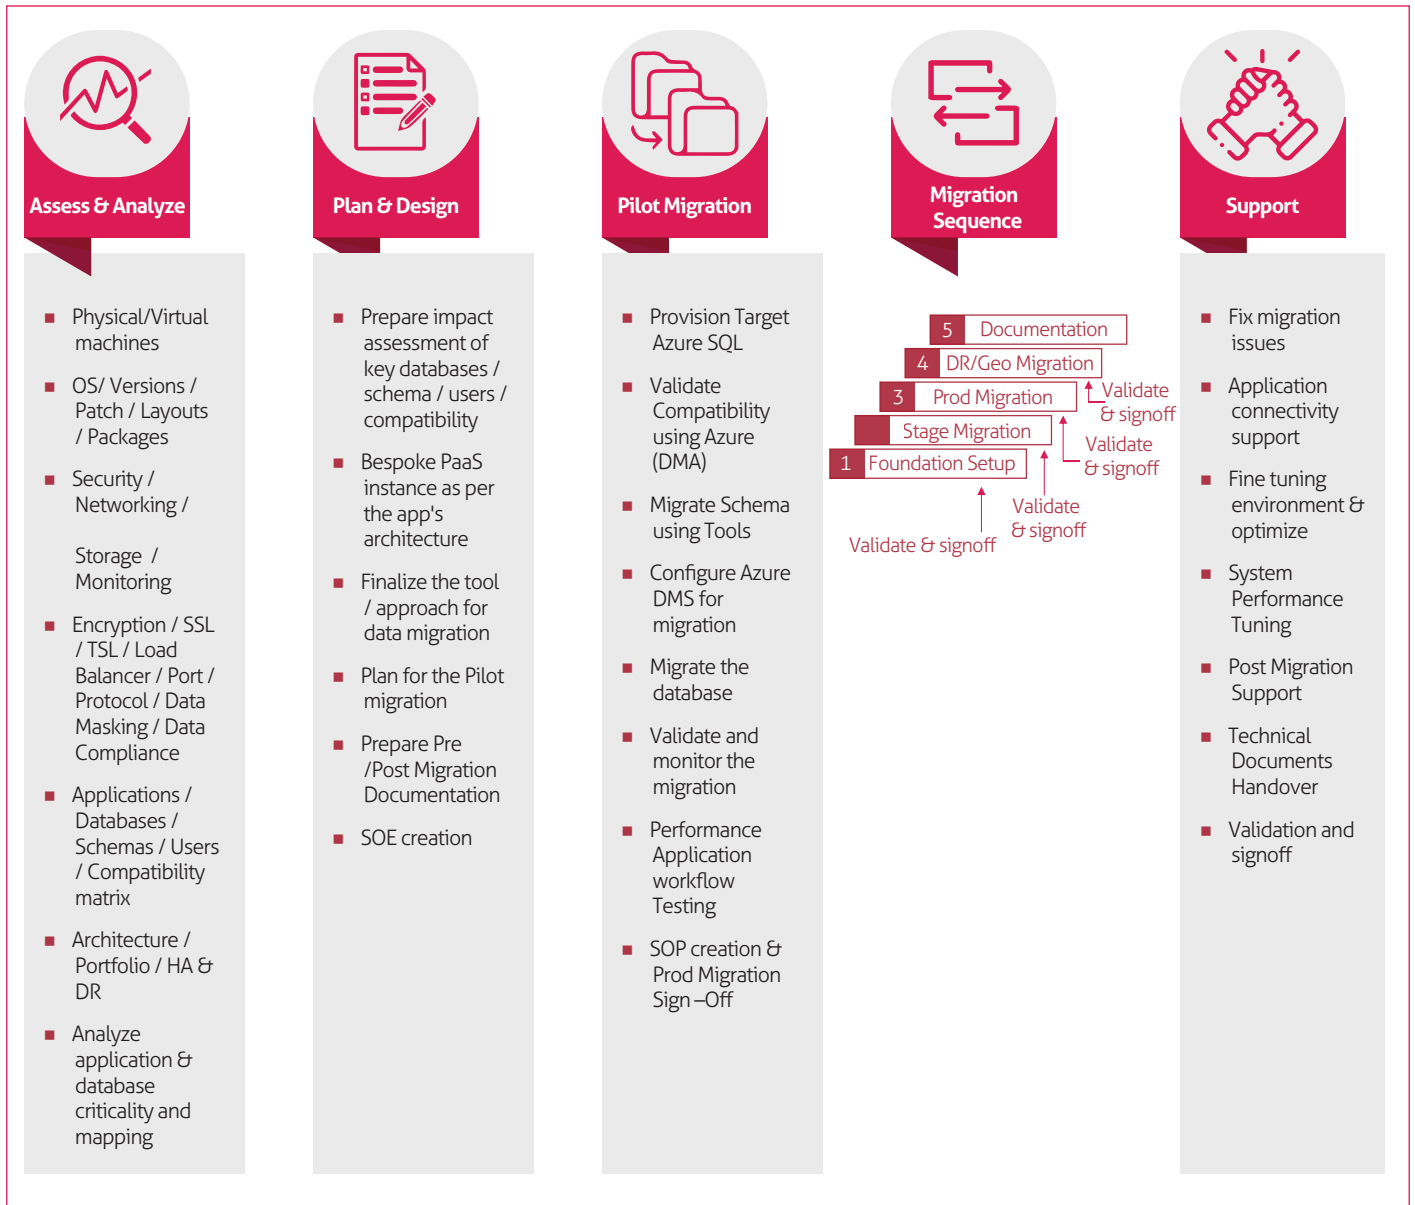
Gone are the days of large-scale database systems that offer more than a thousand functionalities. Customers use less than 20% of these functionalities provided by the specific DBMS/RDBMS systems, and hence, end users are moving towards database systems which offer focused services. It has also been observed that in case an end customer is to integrate more than two-three small-scale database systems for multiple functionalities, it becomes easier and cost effective. Interestingly, most such databases are open source in nature. At the same time, the trend to move things to cloud has also been established as a standard.

In this whitepaper, we have attempted at summarizing a few ways of migrating some of the major open source databases and enterprise databases (like Oracle) to cloud database (predominantly PaaS).



- **Pre-migration:** Analyze the data and create a feasibility report for migration
- **Migration:** Schema Migration and Test Data Migration (Batch Data and Full Load)
- **Post-migration:** Data Reconciliation using Count of Data at Record Level and Column Level, Optimize the tools to ease the process and ensure Ease of Maintenance

# Approach using Migration Framework



In order to support the above-mentioned migration strategy, the framework is divided into five major areas as mentioned below at a high level.

- **Assess and analyze:** Pre-migration Feasibility Report and Analysis
- **Plan and Design:** Migration Plan and Design Documentation, Impact Analysis Roll Back Plan
- **Pilot Migration:** Provision Target Azure SQL DB and Validate Schema and Sample Data
- **Migration Sequence:** Schema, Data, Object, Validate and Verify data
- **Post Migration Support :** Documentation, Support , Fix Bug and Fine Tuning Validation and Signoff

# OSS Use Cases (RDBMS to Azure)









Various use cases for OSS Open Source System are outlined below:

For E.g.: Azure Database for PostgreSQL: Oracle ==> Azure Database for PostgreSQL

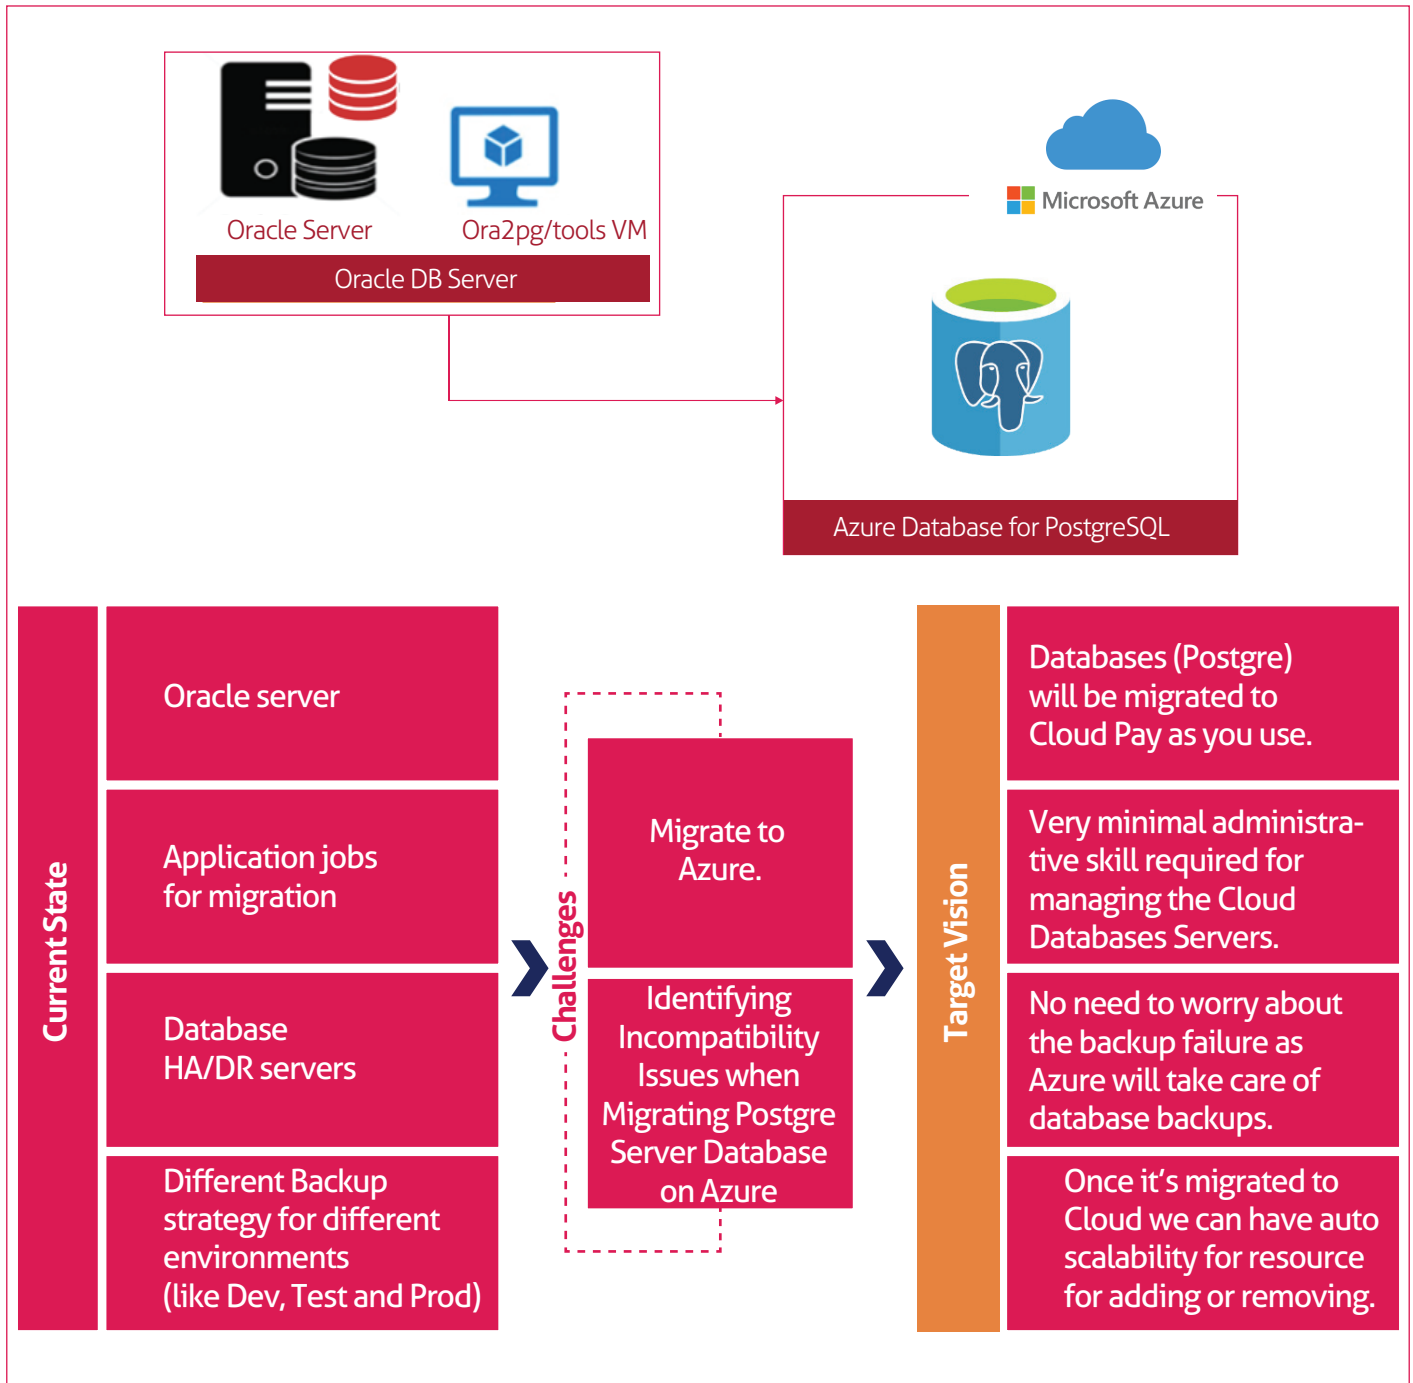
**Target System:** Azure Postgre | **Environment:** Oracle 10g R2, Data Vol 1+ TB, Instances 6+ and Database 5+

**Customer:** Leading airlines in the Middle East | **Tools:** Ora2pg, Flyway, DMA, DMS....etc.

| Use Cases at Mindtree 1/2   |   |   |  |   |
|---|---|---|--|---|
| Source  | Target Azure  | Environment   | Customer   | Tools and Methodology   |
|    |    | <ul style="list-style-type: none"> <li>Oracle 10gR2</li> <li>Data Vol 1TB+</li> <li>Instances 6+</li> <li>Databases 5+</li> </ul>                               | <ul style="list-style-type: none"> <li>Leading A leading Airline in Middle East</li> </ul>                       | <ul style="list-style-type: none"> <li>Ora2pg Flyway DMS</li> </ul>                                 |
|    |    | <ul style="list-style-type: none"> <li>MySQL ver. 5.7x</li> <li>Legacy Systems</li> <li>Data Vol 5TB+</li> <li>Instances 50+</li> <li>Databases 100+</li> </ul> | <ul style="list-style-type: none"> <li>A leading telecommunications and IT service provider in Europe</li> </ul> | <ul style="list-style-type: none"> <li>Workbench</li> <li>DMS</li> <li>DMA</li> <li>SSMA</li> </ul> |
|   |   | <ul style="list-style-type: none"> <li>MySQL 5.6x, 5.7x</li> <li>Data Vol 10TB+</li> <li>Instances 200+</li> <li>Databases 1000+</li> </ul>                     | <ul style="list-style-type: none"> <li>Hi-Tech Media Networks Pvt. Ltd., USA</li> </ul>                          | <ul style="list-style-type: none"> <li>Workbench</li> <li>DMS</li> </ul>                            |
|  |  | <ul style="list-style-type: none"> <li>Oracle 12.1</li> <li>Data Vol 1TB+</li> <li>Instances 2+</li> <li>Databases 2+</li> </ul>                                | <ul style="list-style-type: none"> <li>Insurance Client</li> </ul>   | <ul style="list-style-type: none"> <li>Flyway</li> <li>Ora2pg</li> <li>DMS</li> </ul>               |

| Use Cases at Mindtree 2/2   |   |   |  |   |
|---|---|---|--|---|
| Source  | Target Azure  | Environment   | Customer   | Tools and Methodology   |
|  |  | <ul style="list-style-type: none"> <li>PostgreSQL 9.5, 10</li> <li>Data Vol 1TB+</li> <li>Instances 5+</li> <li>Database 10+</li> </ul> | <ul style="list-style-type: none"> <li>A leading Biotechnology organization</li> </ul> | <ul style="list-style-type: none"> <li>Ora2PG</li> <li>Third party Tools</li> </ul>   |
|  |  | <ul style="list-style-type: none"> <li>MySQL 5.7+</li> <li>Data Vol 500GB</li> <li>Instances 5+</li> <li>Databases 50+</li> </ul>       | <ul style="list-style-type: none"> <li>Global technology Company</li> </ul>            | <ul style="list-style-type: none"> <li>Using logical backup (such as MySQL dump, Mysq pump or and restore is the only option</li> </ul> |
|  |  | <ul style="list-style-type: none"> <li>MongoDB 3.4, 4.0</li> <li>Legacy System</li> <li>Data Vol 2TB+</li> <li>Instances 15+</li> </ul> | <ul style="list-style-type: none"> <li>Global technology Company</li> </ul>            | <ul style="list-style-type: none"> <li>DMA</li> <li>DMS</li> </ul>  |
|  |  | <ul style="list-style-type: none"> <li>SQL Server 2014</li> <li>Data Vol 700GB+</li> <li>Instance 2+</li> <li>Databases 2+</li> </ul>   | <ul style="list-style-type: none"> <li>A leading brand in FMCG in India</li> </ul>     | <ul style="list-style-type: none"> <li>SSMA</li> <li>DMA</li> </ul>   |

# PoC: Oracle to Azure PostgreSQL Using Compass and Ora2pg



PostgreSQL is open source, and has become the choice for several enterprise developers. We will be able to deploy scalable PostgreSQL deployments in minutes on Azure with a cost-effective solution.

As part of this exercise, we have classified the solution into the following sections:

- **Current State:** Oracle Server, Database, Object, Applications for Migration
- **Challenges:** Schema and Object (PLSQL Code) Migration, Data Challenges
- **Target Version:** End State Flexible, Scalable and Durable Solution

# Recommended Tools for Migration (Compass)

### Your Infrastructure

Information gathered by ComPass® after scanning your infrastructure.

| Scan Summary          |                  |
|-----------------------|------------------|
| No. Of Servers        | 60               |
| Windows Servers       | 60               |
| Linux Servers         | 0                |
| Unix Servers          | 0                |
| Workload Components   | 2389             |
| Physical Servers      | 5                |
| Virtual Servers       | 55               |
| Total Storage         | 11.4375 TB       |
| No. Of Disks          | 170              |
| No. Of Drives Size    | 178 [10.9336 TB] |
| Unpartition Disk Size | [516 GB]         |

### Memory Utilization

Memory utilization of your application server at the time of scanning.

### Workload Components Distribution

Major workload components such as web, database, mail, queue, LDAP servers, etc.

### Your Preferred Clouds

| Your Preferred Clouds | Cloudability Status |
|-----------------------|---------------------|
| AWS                   | Cloudable           |
| Azure                 | Cloudable           |
| GCP                   | Cloudable           |

### PaaS TCO Match Over 3 Years

| Cloud Name      | Total Recurring /Month | Total Cost Over 3 Years |
|-----------------|------------------------|-------------------------|
| Microsoft Azure | \$691                  | \$24876                 |
| AWS             | \$388                  | \$13968                 |

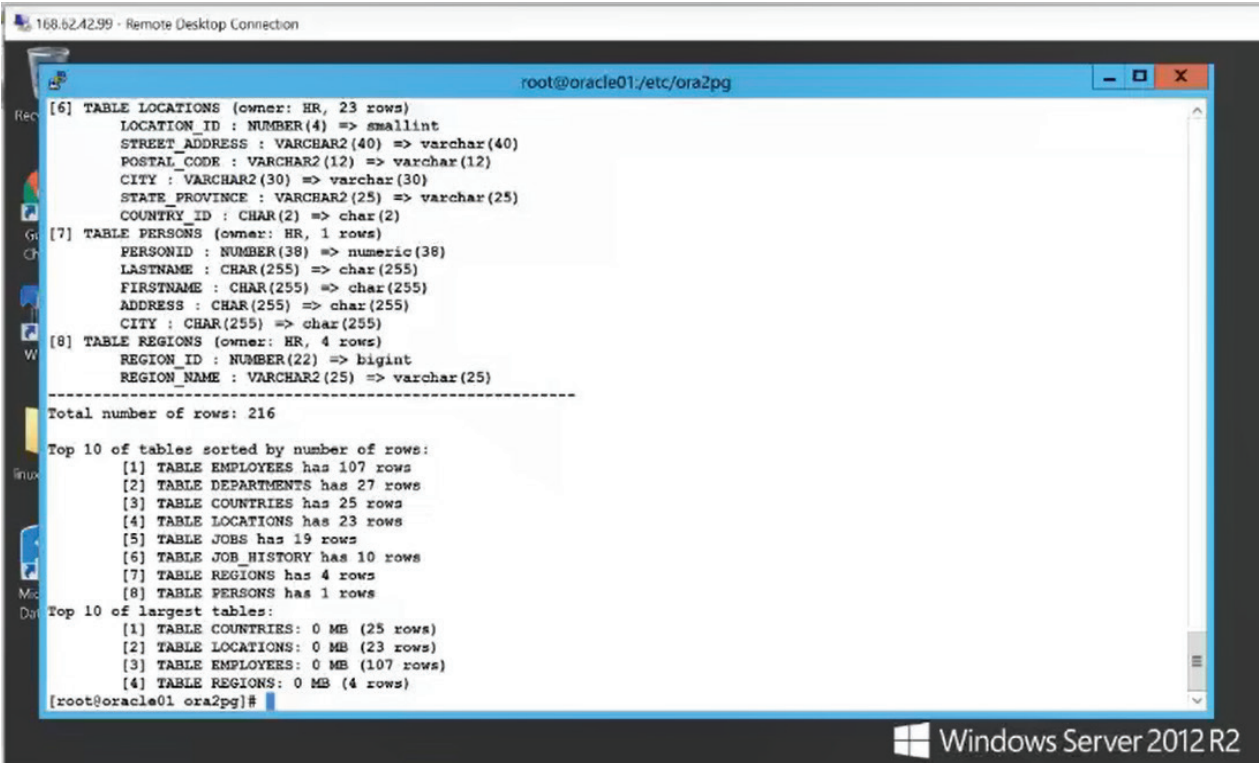
## About Compass:

- Compass is Mindtree's accelerator to discover, analyze, assess, plan and mitigate risks
- It enables accelerated cloud adoption by performing and orchestrating all the activities
- Extended Performance Monitoring (EPM) creates performance reports for the server
- The infrastructure profile provides you with an overview of the datacenter
- Creates application profiles to modernize your application
- Servers for particular applications and create application profile
- Performs scanning in agentless and agent-based approach
- Groups all the servers in the cluster and creates an infrastructure profile
- Provides suggestions for storage advisory and to modernize your application
- Delivers PaaS recommendation for the workload which can be migrated AS-IS



# Recommended Tools for Migration (Ora2pg)

Ora2pg helped ease the process of migration. Here, we migrated the sample HR database from an on-premises Oracle 11g to an Azure PostgreSQL database by using the Ora2pg and Azure Database Migration Service.



```
root@oracle01:/etc/ora2pg
[6] TABLE LOCATIONS (owner: HR, 23 rows)
LOCATION_ID : NUMBER(4) => smallint
STREET_ADDRESS : VARCHAR2(40) => varchar(40)
POSTAL_CODE : VARCHAR2(12) => varchar(12)
CITY : VARCHAR2(30) => varchar(30)
STATE_PROVINCE : VARCHAR2(25) => varchar(25)
COUNTRY_ID : CHAR(2) => char(2)
[7] TABLE PERSONS (owner: HR, 1 rows)
PERSONID : NUMBER(38) => numeric(38)
LASTNAME : CHAR(255) => char(255)
FIRSTNAME : CHAR(255) => char(255)
ADDRESS : CHAR(255) => char(255)
CITY : CHAR(255) => char(255)
[8] TABLE REGIONS (owner: HR, 4 rows)
REGION ID : NUMBER(22) => bigint
REGION_NAME : VARCHAR2(25) => varchar(25)
-----
Total number of rows: 216
Top 10 of tables sorted by number of rows:
[1] TABLE EMPLOYEES has 107 rows
[2] TABLE DEPARTMENTS has 27 rows
[3] TABLE COUNTRIES has 25 rows
[4] TABLE LOCATIONS has 23 rows
[5] TABLE JOBS has 19 rows
[6] TABLE JOB_HISTORY has 10 rows
[7] TABLE REGIONS has 4 rows
[8] TABLE PERSONS has 1 rows
Top 10 of largest tables:
[1] TABLE COUNTRIES: 0 MB (25 rows)
[2] TABLE LOCATIONS: 0 MB (23 rows)
[3] TABLE EMPLOYEES: 0 MB (107 rows)
[4] TABLE REGIONS: 0 MB (4 rows)
[root@oracle01 ora2pg]#
```

## Here are a few key technical benefits of using Azure PostgreSQL:

- Database migration assessment from Oracle to Postgre
- Ease in creation of migration projects and plan
- Automated discovery using Oracle Database Discovery
- Automation (full and delta) in database schema export
- Ease in automatic conversion of PL/SQL to PLPGSQL
- Allow autonomous transaction


### Benefits:

- Below are long-term benefits in terms of the demo performed for Oracle to Azure PostgreSQL Reduction in cost (Open Source Licenses)
- No licenses fees (No Fee for the Community Edition on Postgre)
- Scalable on demand on Azure
- Flexible, fast and predictable storage
- Reduction in efforts, monitoring and metrics
- Global reach, high availability
- Ease of conversion, isolation and security
- Reduction in operation cost due to easy and managed deployments

# Contact details:

Saptarshi Biswas - [Clouddatappractice@mindtree.com](mailto:Clouddatappractice@mindtree.com)

*This document is the exclusive property of Mindtree Limited (Mindtree). The recipient agrees that they will not copy, transmit, use or disclose the confidential and proprietary information in this document by any means without the expressed and written consent of Mindtree. By accepting a copy, the recipient agrees to adhere to these conditions to the confidentiality of Mindtree's practices and procedures; and to use these documents solely for responding to Mindtree's operations methodology.*



## Welcome to possible

### About Mindtree

Mindtree [NSE: MINDTREE] is a global technology consulting and services company, helping enterprises marry scale with agility to achieve competitive advantage. "Born digital," in 1999 and now a Larsen & Toubro Group Company, Mindtree applies its deep domain knowledge to 280+ enterprise client engagements to break down silos, make sense of digital complexity and bring new initiatives to market faster. We enable IT to move at the speed of business, leveraging emerging technologies and the efficiencies of Continuous Delivery to spur business innovation. Operating in more than 15 countries across the world, we're consistently regarded as one of the best places to work, embodied every day by our winning culture made up of over 21,800 entrepreneurial, collaborative and dedicated "Mindtree Minds".

To learn more about us, visit [www.mindtree.com](http://www.mindtree.com) or follow us @Mindtree\_Ltd