

WHITE PAPER

ORACLE MIGRATION TO POSTGRES

An end-to-end methodology
to Achieve Your Business Goals.

Marc Linster, EDB Technical Fellow
Raghavendra Rao, Global Migration, Practice Leader



Contents

1	Introduction	3
2	Defining the migration scope and the ROI	4
2.1	Understanding the database estate	4
2.2	Understanding the business case and the ROI	5
2.3	Creating the Prioritized Migration Backlog	8
2.4	Migration Sprints	8
3	Migration and Compatibility Considerations	9
3.1	Do it yourself or work with EDB	9
3.2	Training and enablement	9
3.3	Databases don't exist in isolation	9
4	Summary	9

1 INTRODUCTION

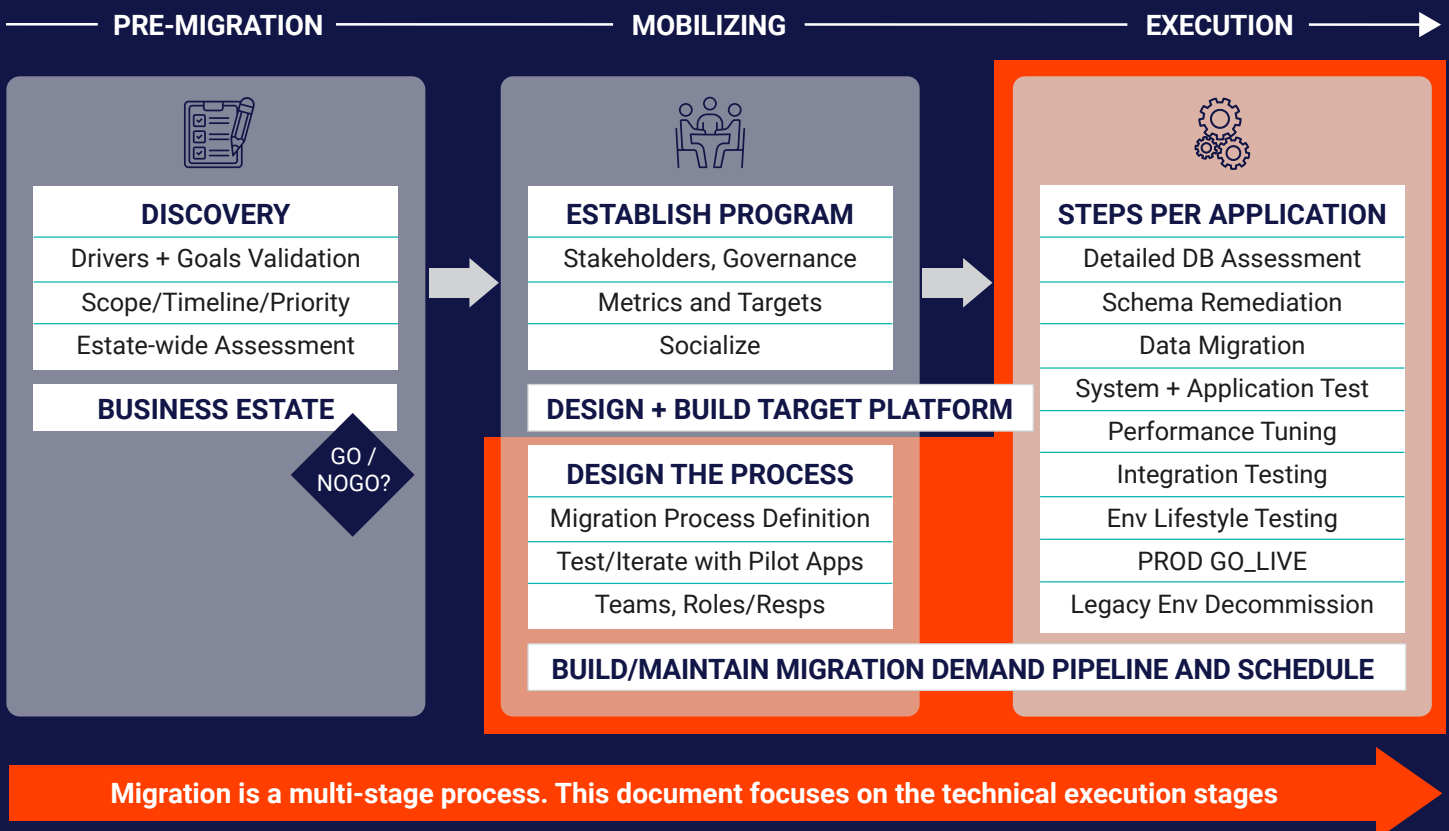
Migrations off of legacy databases, especially off of Oracle databases, have become a core component of digital transformations. There are multiple business drivers for these migrations, ranging from cost reduction to agility, innovation, and the move to the cloud. Traditionally, customers asked EDB to migrate individual application databases. Today we are increasingly encountering wholesale strategic moves onto open source-based databases. “We have 1000 databases, how can you help us?” has become a standard question asked as part of wholesale digital transformations.

In this whitepaper we are describing EDB’s methodology to address these large scale migration challenges. We will

walk the reader through all the technical planning steps, starting with the assessment of the database estate, identifying Postgres migration candidates, assessing them for compatibility and ease of migration, defining the business case and the ROI, and executing the migrations through a series of sprints.

Moving from one database technology to another is not only a technological exercise. It also includes mobilizing the team, building the skills, defining the projects, creating a governance model, and significant change management. Those aspects are described elsewhere; this paper discusses the technology aspects only.

Figure 1: Database migration change management steps



2 DEFINING THE MIGRATION SCOPE AND THE ROI

EDB’s methodology takes an iterative approach. In a first step we assess the current database estate, identify which databases should be part of the scope, and we define the ROI for the database migration strategy. This defines the Prioritized Migration Backlog.

We then iterate through the backlog in multiple sprints - each sprint focused on several databases (the Sprint Backlog) that will be migrated to EDB Postgres and moved to production.

After every sprint, we review the Prioritized Migration Backlog, pick a new set of databases, and re-execute the migration and move to production.

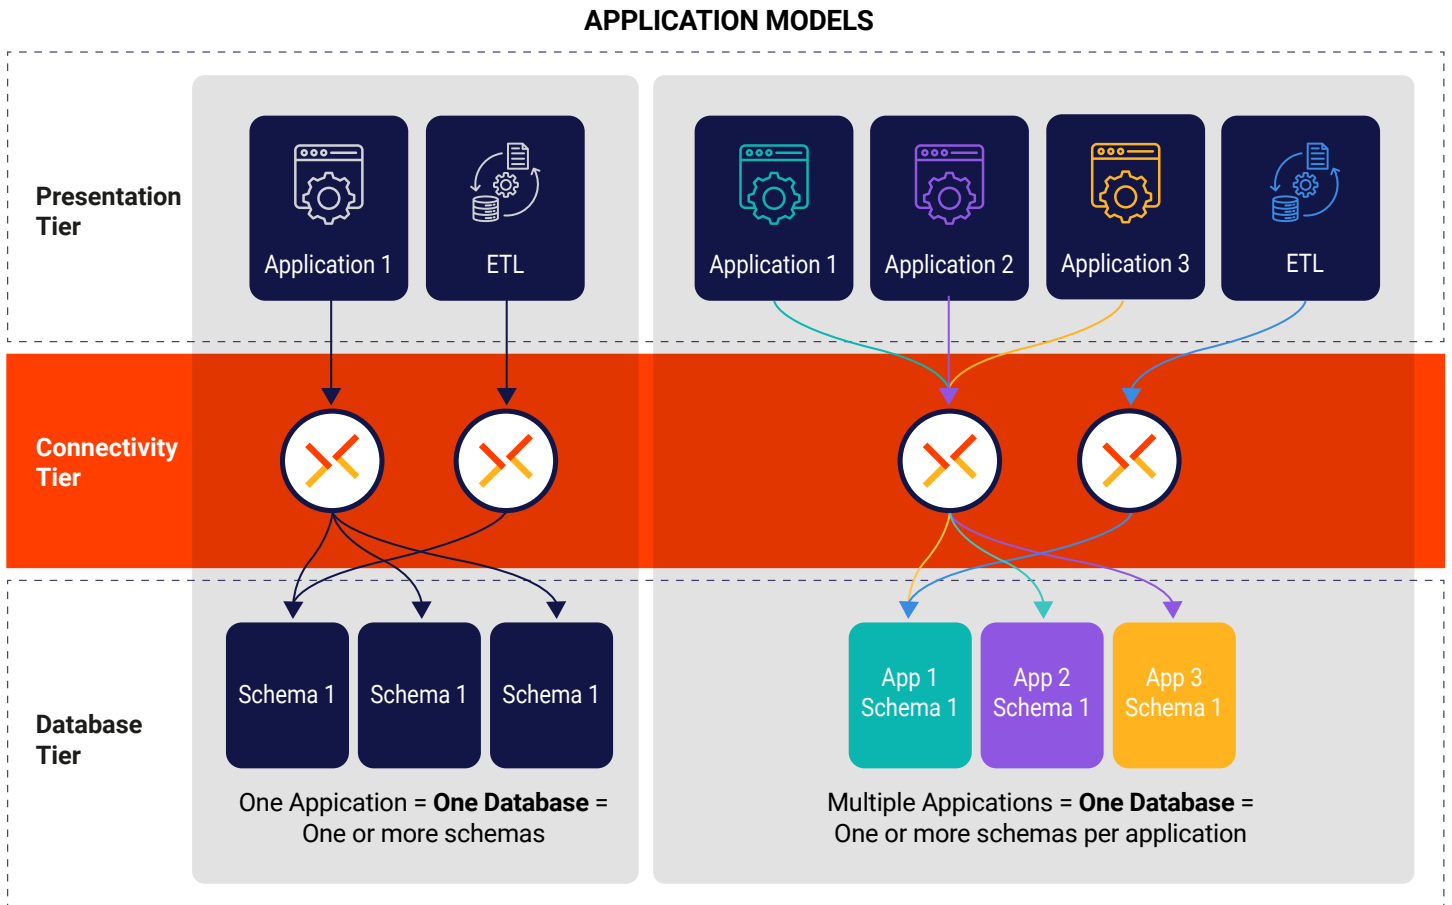
If desired, advanced Scrum concepts, such as velocity, product owner, Scrum master, etc. can be integrated with this methodology.

2.1 Understanding the database estate

First and foremost, we have to create a list of the enterprise’s databases and understand the database estate, how it links to the application landscape, and where the applications are in their lifecycle. Apps that will be retired soon, or applications that are due for a major technical overhaul are most often excluded from a migration plan.

Understanding the link between applications and databases is key, as applications tend to use more than one database, and databases can be connected to several applications.

Figure 2: Linking applications and databases



Data Lake	Data Warehouse	Data Mart	ODS	System of Record	System of Engagement	Edge Computing	Mobile Computing
Description 100 TB+; read-only; relational +; multiple topics; rel. Integrity not relevant ; data scientist focus	50TB+; read only; relational; star/cube/mat. view; automatic refresh mat. view; multiple topics; analyst focused	10TB+; read only; relational; star/cube/mat. view; periodic refresh mat view; single topic; citizen analyst focus	10TB+; read only; relational only; normalized data model; concurrent refresh; citizen analyst focused	1-20TB; read/write; 90% relational; ACID; ERP/ CRM, Line of business focused	< 2TB; read/write; relational + doc + GIS; ACID or BASE; micro services; website and social media focus	< 2TB; read/write; relational + doc + GIS; ACID or BASE; micro services; M2M, IoT focus	< 1GB; read/write; relational + doc + GIS; single user; micro services IoT, mobile app focus
Key players Hadoop, Cassandra, Data Bricks	Redshift, Snowflake, Teradata, Exadata, BigQuery	Exadata, Oracle DB, SQL Server, Postgres	Oracle, SQL Server, Postgres	Oracle, SQL Server, Postgres	MongoDB, MySQL, Redis, Postgres	MySQL, Postgres	SQLite, Realm
Postgres strengths		Analytics capab.; compatibility; ease of integration	Analytics capab.; compatibility; ease of integration	Innovation; cost; compatibility; no vendor lock in;	Multi-model; innovation; cost; Available everywhere	Multi-model; innovation; cost; Available everywhere	
# in the enterprise	1 (maybe 2-5)	10-50	5-20	10-5,000	10-2,000	100 -10,000	10,000+

Table 1: Areas of primary applicability for Postgres

Today, core Postgres is most often a good fit for data marts, operational data stores, systems of record, systems of engagement, and edge solutions. While Postgres has been used for data lakes, data warehouses and for mobile applications, it is not generally considered as a best practice platform for these use cases.

For this initial assessment, we recommend to collect the following information

- Application name
- Commercial off the shelf (COTS) or custom developed (Y/N)
- Nbr of database(s)/schema(s)
- Primary use case (data lake, data warehouse, data mart, ODS, SoR, SoA, Edge, Mobile)
- Data volume and read/write % balance
- Lifecycle stage (strategic, EoL within 24 months, scheduled for overhaul within 24 months)
- Current database technology platform (Oracle, DB2, SQL Server, Sybase, ...)
- Is the application development or maintenance team available for verification, and potentially, minor changes to the application

Using the criteria outlined above, you can use data in Table 1 to quickly assess which part of the estate can be targeted for migration to Postgres.

2.2 Understanding the business case and the ROI

The applications that are targeted to go to Postgres will undergo an initial compatibility assessment with EDB Postgres Advanced Server (EPAS). EPAS is EDB's Oracle-compatible distribution of Postgres that significantly simplifies the migration to Postgres, as EPAS natively understands most Oracle database programming constructs and provides drivers that are largely compatible with the Oracle database drivers.

The initial high level analysis puts the application databases into easy, medium, hard or custom categories. EDB's experience allows us to estimate roughly how difficult those migrations will be, how much effort is required, and what the cost will be to execute the migrations. By analyzing the core count and add-on features that were licensed from Oracle, we can estimate the cost savings that would result from the migration. The combination of the cost to execute the migrations and the license savings, will provide enough information to define the Return on Investment (ROI) and a high-level timeline for the migration.

To assess the degree of compatibility and the migration effort, EDB uses the EDB Compatibility Assessment Tool (ECAT). ECAT connects to every database to extract information about the database schema, the stored procedures, the data volumes, the number of database cores and the key Oracle add-ons like RAC, Data Guard, etc. ECAT can be operated with read-only access to the Oracle database catalogs. Additionally, the analysis requires filling in a short questionnaire about the application and the database driver.

The questions are:

- What programming language is being used for the application?
- What driver is used to connect the application to the database?
- Is the application team available for verification testing and potentially necessary modifications?
- How many database environments (e.g., dev, performance test, production staging, DR replicas, etc) have to be instantiated as part of the migration?
- What are the HA requirements (99%, 99.5%, 99.9%, or 99.99%) and geo-redundancy requirements, and how are they met today?

ECAT generates high-level assessment reports and estimates the Compatibility Class (very easy, easy, medium, hard, custom).

The Compatibility Classes are defined as follows

- **Very Easy:** No conversion at all/No incompatible code objects
- **Easy:** Minor changes/Automatic repair through Migration Portal or known, simple workarounds.
- **Medium:** Moderate conversions need to be applied.
- **Hard:** New design, development and testing needed
- **Very Hard:** Complete rewrite of application, db features etc (Eg: Apex, Forms) that may not be worthwhile targets for a migration effort

Each summary line is supported with a detailed explanation

- Detailed Report showing the output of the Compatibility Assessment produced for each application.
- The report lists the names of the incompatible Oracle features discovered in the analysis.
- For each incompatible feature, it highlights the number of occurrences of each feature encountered (used within the estimation calcs).

Table 2: SampleECAT summary report with explanation

High Level Compatibility Estimates:-							
Application_Name	Schemas	Features	Occurrences	Avg. Schema Conversion Hours	Max Score	Compatability_Class	Total Migration Days
raptor	1	1	0	4	1	Very Easy	1
raptor_snap	1	4	16	32	3	Medium	4
raptor_sync	9	12	80	180	4	Hard	23

Compatibility Class

- **Very Easy:** No conversion at all/No incompatible code objects
- **Easy:** Minor changes/Automatic repair through Migration Portal or known, simple workarounds.
- **Medium:** Moderate conversion needs to be applied.
- **Hard:** New design, development and testing needed
- **Very Hard:** Complete rewrite of application, db features etc (Eg: Apex, Forms)

Databases, Schemas, Features, Occurrences

- Counts by application of databases, schemas, features identified and the number of occurrences of the features.

Average Conversion/ Migration Days

- This is an estimate of the average effort required to convert the schema and migrate the data based on the high-level compatibility assessment.
- From EDB's experience this estimate is within +/- 20% in 80% of scenarios.

	Application Name	Schemas	Incompatible Features	Occurrences
0	raptor	all_schemas	No	0
1	raptor_snap	license	dbms_lob	4
2	raptor_snap	license	dbms_sqlltune	3
3	raptor_snap	license	utl_http	4
4	raptor_snap	license	utl_i18n	5
5	raptor_sync	app_password_owner	dbms_assert	43
6	raptor_sync	ard_amer_owner	dbms_stats	1
7	raptor_sync	ard_amer_user	dbms_application_info	1
8	raptor_sync	ard_amer_user	dbms_ddl	1
9	raptor_sync	license	dbms_lob	4
10	raptor_sync	license	dbms_sqlltune	3
11	raptor_sync	license	utl_http	4
12	raptor_sync	license	utl_i18n	5
13	raptor_sync	rbsdba_jgos	dbms_shared_pool	1
14	raptor_sync	rcm_user	dbms_application_info	1
15	raptor_sync	rcm_user	dbms_ddl	1
16	raptor_sync	sam_app_user	dbms_application_info	1
17	raptor_sync	sam_app_user	dbms_ddl	1
18	raptor_sync	sam_prof_user	dbms_application_info	1
19	raptor_sync	sam_prof_user	dbms_ddl	1
20	raptor_sync	stg_owner	dbms_application_info	3
21	raptor_sync	stg_owner	dbms_stats	2
22	raptor_sync	stg_owner	dbms_xmlgen	3
23	raptor_sync	stg_owner	owa_util	2

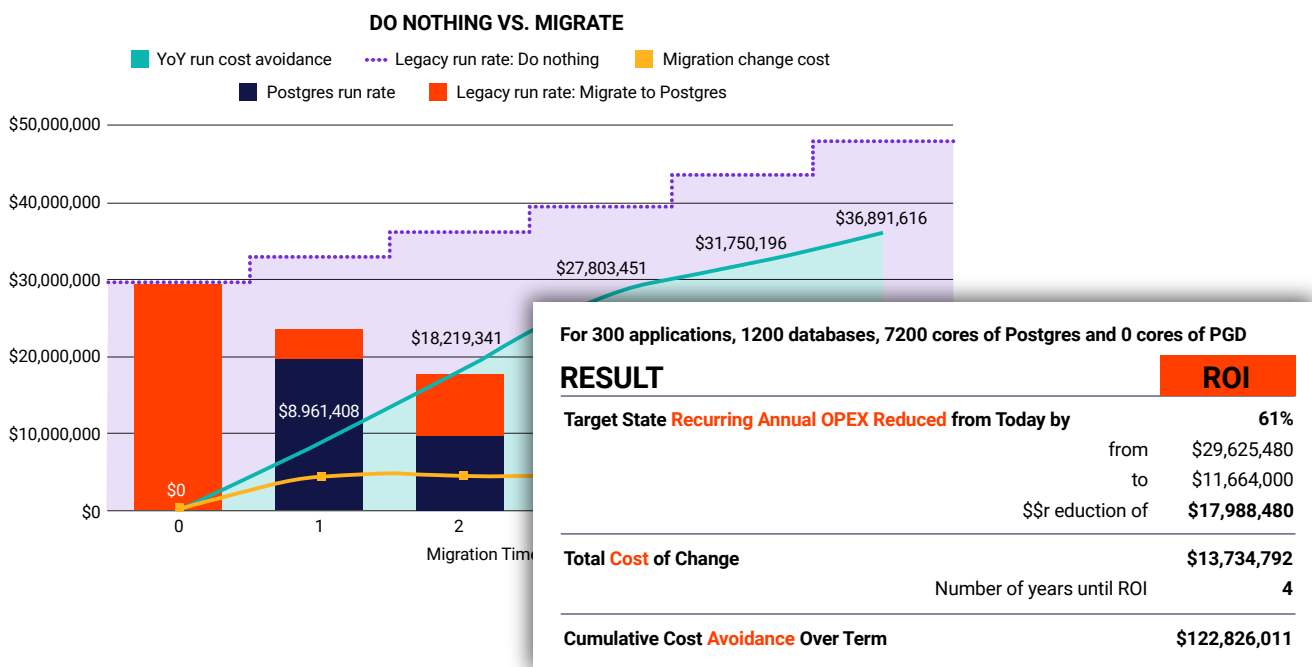
Figure 3: Detailed compatibility report

While the EDB migration methodology leverages the native Oracle compatibility of EDB Postgres Advanced Server, a subset of the applications in Compatibility Class 'Very Easy' can also be migrated to PostgreSQL, the open source distribution of Postgres that does not include Oracle compatibility. If desired, an additional step can be performed for the 'Very Easy' applications using ORA2PG.

If ORA2PG classifies an application as 'Migration Level A/ Technical Level 3', then the database is likely to run on PostgreSQL and does not need the additional capabilities provided by EPAS.

The Compatibility Class definition and high level effort assessment feed into EDB's ROI Estimator for Oracle Migrations.

Figure 4: Sample output of the EDB ROI Estimator for Oracle Migrations



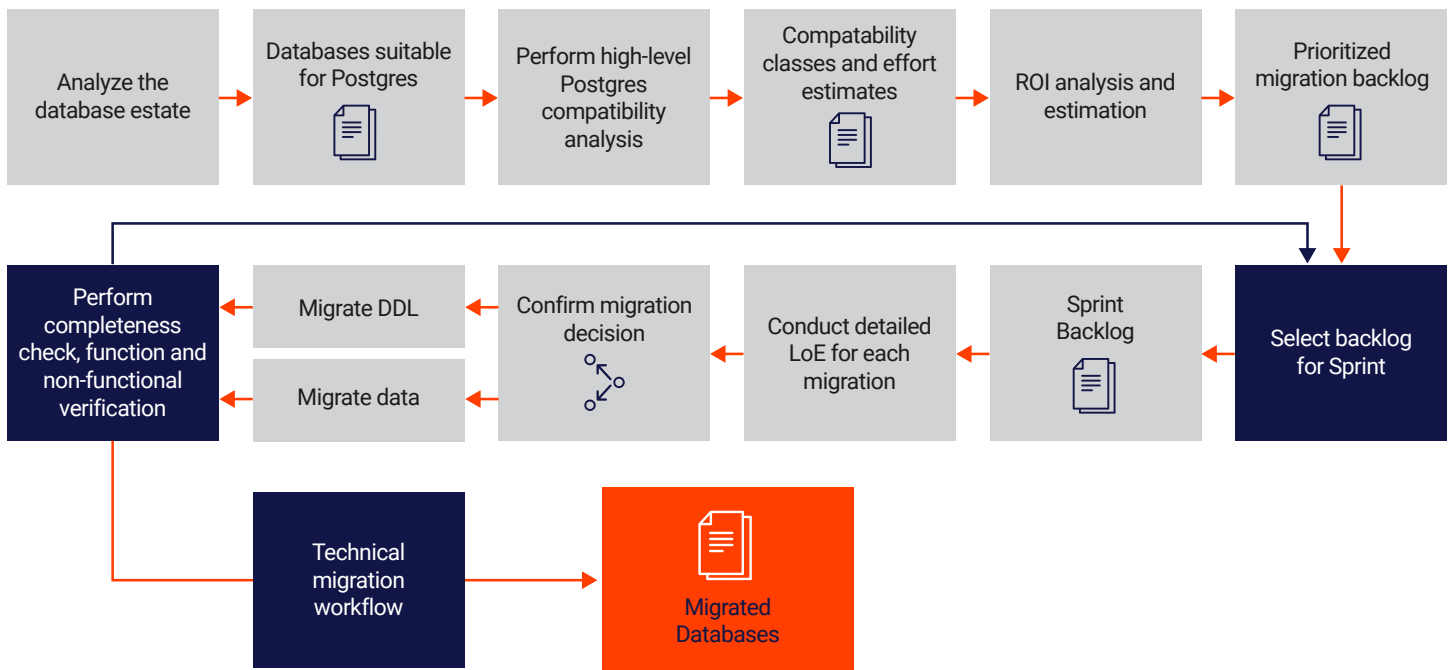


Figure 5: Technical Migration Workflow

The ROI Estimator requires information about actual license cost, estimated discounts, labor rates, and project duration.

The ROI and the ECAT report will be the basis for the strategic decision to migrate and to define the scope of the databases to be migrated, aka the Migration Backlog. The backlog includes the application databases that are considered for migration.

The Migration Backlog includes the following

- Application Name
- Database(s)/Schemas
- High Level compatibility classification (very easy, easy, medium, hard, very hard)
- High level effort and cost estimate
- License cost reduction estimate

Database migration is an iterative process that starts with identifying the databases that should be considered, gaining a high-level understanding for the feasibility and the ROI, before diving into the actual migration process.

2.3 Creating the Prioritized Migration Backlog

The Migration Backlog will be prioritized based on the business case, the business priorities, application

considerations, infrastructure requirements, and technical skill set availability.

We highly recommend starting migration projects with applications in the Compatibility Class 'Very Easy' or 'Easy' that do not have significant HA requirements, e.g., that were not migrated from RAC or Golden Gate, but that still have meaningful core counts. Availability of the application team to support functional and non-functional verification is a key factor too. Prioritizing 'Very Easy' and 'Easy' applications will give the migration team the opportunity to build up their processes and skill sets, before tackling more difficult challenges.

This step creates the Prioritized Migration Backlog.

2.4 Migration Sprints

During the sprints, the team picks a small number of applications off the top of the Prioritized Migration Backlog. Depending on the size and experience of the team, we encourage you to start with 1-5 applications.

For each application we recommend the following steps

Create a detailed Level of Effort (LoE) for the migration of each application. The ECAT tool provided an estimated classification into Very Easy/Easy/Medium/Hard/Very Hard based on a fully automated rapid assessment suitable to understand a large database estate. This classification is within 20% of accuracy for the majority of migrations.

The detailed LoE uses EDB's Migration Portal to review the DDL and the stored procedures in detail and identify the work to be done to complete the migration. EDB also uses EDB's SQL Assess to extract application-generated SQL code and assess it for EDB Postgres compatibility. We may also request additional information, such as Oracle Automatic Workload Repository (AWR) reports, to complete the detailed LoE.

The data needed for the detailed LoE has already been collected in ECAT. The detailed LoE is produced by EDB's Migration Factory, a team specialized on rapid and cost effective Oracle migrations.

Besides the code migration, the estimate should include creating the architecture, deploying the infrastructure,

integration with enterprise environments, converting the schema, migrating the data, validating the data, and providing support for application functional testing.

The migration of the schema (DDL and stored procedures) can happen in parallel with the setup of the infrastructure.

Data migration can be done as a snapshot (we estimate approx 50GB/Hour within the same data center) using EDB's Migration Toolkit (MTK), but that depends on the infrastructure and the latency. Consider EDB Postgres Replication Server if change data capture is required to minimize down time.

The completeness of the data migration should be verified using EDB's LiveCompare.

3 MIGRATION AND COMPATIBILITY CONSIDERATIONS

3.1 Do it yourself or work with EDB

EDB's Professional Services team, and EDB's Migration Factory, have done Oracle migrations for almost 20 years. While we make the majority of our tools available to our customers and prospects, not all of our know-how has been codified. There are also several tools, such as ECAT or the SQL Assess tool, that are currently only available in conjunction with an EDB professional services engagement.

We encourage customers to create blended teams that include customer staff or the customer's preferred systems integrator, so that the EDB team can focus on teaching and enabling the customer's team

3.2 Training and Enablement

EDB has made most of our training classes available for free to help customers and prospects accelerate their

migration to Postgres (<https://www.enterprisedb.com/training/on-demand>). We highly recommend that the team take the EDB Postgres Advanced Server courses, including classes for developers and classes focused on performance tuning.

EDB can also work with you to develop customized learning and certification paths.

3.3 Databases don't exist in isolation

While database migration projects focus on the database, they often require support from the application teams. EDB Postgres is highly Oracle compatible, but even if the migration does not require any application changes, the functional and non-functional verification activities require participation from the application team.

3 SUMMARY

Migration off legacy databases, especially off Oracle, is a key component of digital transformation. It is fundamental to leveraging commodity hardware, establishing cloud choice, and enabling the move to cloud-native

microservices architectures. Migration can be a daunting task, but with the right tools and the right process one can mitigate the risk and execute quick and painless migrations in a predictable fashion.



ABOUT EDB

EDB provides enterprise-class software and services that enable businesses and governments to harness the full power of Postgres, the world's leading open source database. With offices worldwide, EDB serves more than 1,500 customers, including leading financial services, government, media and communications and information technology organizations. As one of the leading contributors to the vibrant and fast-growing Postgres community, EDB is committed to driving technology innovation. With deep database expertise, EDB ensures extreme high availability, reliability, security, 24x7 global support and advanced professional services, both on premises and in the cloud. This empowers enterprises to control risk, manage costs and scale efficiently. For more information, visit www.enterprisedb.com.