

Aurora PostgreSQL Artifact Extraction Prerequisites

Contents

- 1. Introduction3
- 2. Artifact Extraction.....3
 - 2.1 DDL Scripts3
 - 2.2 Query Execution Logs3
 - 2.2.1 Config File Location3
 - 2.2.2 Data Directory Path.....3
 - 2.2.3 Logging4
 - 2.2.4 Restarting Services.....4
 - 2.2.5 Download Log Files4
 - 2.3 Other Database Objects5
 - 2.4 Other Scripts and Artifacts6
- 3. Getting Help.....7

1. Introduction

LeapLogic's Assessment profiles existing inventory, identifies complexity, performs dependency analysis, and provides recommendations for migration to modern data platform.

2. Artifact Extraction

LeapLogic requires certain artifacts to perform an assessment. As a prerequisite, to fetch the required data you must have admin privileges. You can retrieve these artifacts from your GIT instance or the PostgreSQL repository where all the artifacts such as DDL scripts, stored procedures, functions, query execution logs, DML scripts, and other database objects are stored. LeapLogic needs all these artifacts in the form of .sql files.

2.1 DDL Scripts

One of the easiest and most direct ways to export the required DDLs from the database is by using the `pg_dump` command in PostgreSQL.

Note: It is recommended to execute the `pg_dump` command for all databases, particularly the ones you intend to migrate.

Use the following command to export the DDL scripts to the machine where PostgreSQL is installed. Modify the command details highlighted in **Yellow** based on your requirements. This command will export all the DDLs of the database objects present in the PostgreSQL database.

```
pg_dump -h host -U username -d database -p 5432 --schema-only -f /path/pg_dump_database_name.sql
```

2.2 Query Execution Logs

To export the query execution logs, make sure that the error reporting and logging are enabled in PostgreSQL. Follow the below given steps to enable logging.

2.2.1 Config File Location

Run the command below in PostgreSQL to find the location of the configuration file.

```
postgres=# SHOW config_file;
```

2.2.2 Data Directory Path

The data directory is the storage location for the query execution logs. Run the command given below in PostgreSQL to find the data directory path.

```
postgres=# SHOW data_directory;
```

2.2.3 Logging

1. Open the Amazon RDS console and choose **Databases** from the navigation menu.
2. Choose the DB instance for which you want the logging.
3. Choose the **Configuration** view to see the details of the DB instance and to see the parameter group associated with the DB instance.

Note: When you create a DB instance, it is associated with the default DB parameter group. Because you can't modify this group, create a custom parameter group to modify it.

4. Choose **Parameter group** from the **Instance configuration** pane, and then choose the parameter group that you want to modify.
5. Click **Edit Parameter**.
6. In the **Filter parameters** field, select the parameters given below to change:

Enter **log_statement** and change the value to **all**.

Enter **log_destination** and change the value to **csvlog**.

Enter **log_filename** and change the value to **postgresql.log.%Y-%m-%d-%H%M**.

7. Click **Save changes**.

2.2.4 Restarting Services

When you change the DB instance parameter group (from default to a custom group), you must reboot the DB instance.

Copy the log file from the data directory location and share it with the LeapLogic team.

2.2.5 Download Log Files

Follow the steps below to download the log files from the AWS Management Console.

1. Open **Amazon RDS**.
2. Choose your **RDS/Aurora instance**.
3. Choose **Logs and events**.
4. In Logs, choose the **all-csv** log files.
5. Click **Download**.

The top screenshot shows the Amazon RDS console for the 'll-aurora' database instance. The 'Logs & events' tab is selected, and the 'Auto scaling policies (0)' section is visible below it.

The bottom screenshot shows the 'Logs (48)' section. A list of log files is displayed, including 'error/postgresql.log.2024-10-23-0200' and 'error/postgresql.log.2024-10-23-0500.csv'. The 'Download' button is highlighted.

Share all the CSV log files with the LeapLogic team.

2.3 Other Database Objects

To enhance the assessment results for your environment and workloads, it is recommended to export additional database objects. Please refer to the below script to export the required database objects as separate delimited files.

```
--- 2.3.1 Database Objects
SELECT table_name,
table_catalog AS databasename,
table_schema AS SchemaName,
table_type
FROM information_schema.tables
WHERE table_type = 'BASE TABLE'
AND table_schema NOT IN
('pg_catalog', 'information_schema')
UNION
SELECT table_name,
```

```

table_catalog AS databasename,
table_schema AS SchemaName,
'VIEW' AS table_type
  FROM information_schema.views
 WHERE table_schema NOT IN ('pg_catalog', 'information_schema')
   AND table_name !~ '^pg_';

--- 2.3.2 Database Objects Count.
SELECT databasename,
table_type,
count(table_name) AS table_count
  FROM (SELECT table_name,
table_catalog AS databasename,
table_schema AS SchemaName,
table_type
  FROM information_schema.tables
 WHERE table_type = 'BASE TABLE'
   AND table_schema NOT IN
     ('pg_catalog', 'information_schema'))
UNION
SELECT table_name,
table_catalog AS databasename,
table_schema AS SchemaName,
'VIEW' AS table_type
  FROM information_schema.views
 WHERE table_schema NOT IN ('pg_catalog', 'information_schema')
   AND table_name !~ '^pg_') as subquery
GROUP BY 1, 2;

--- 2.3.3 Database Volume
select t1.datname AS db_name,
       pg_size_pretty(pg_database_size(t1.datname)) as db_size
from pg_database t1
order by pg_database_size(t1.datname) desc;

--- 2.3.4 Table volume
SELECT
table_catalog AS databaseName,
table_schema,
table_name,
pg_size_pretty(pg_total_relation_size(quote_ident(table_schema)||'.'||quote_ident(table_name))) AS total_size
from information_schema.tables
where table_schema = 'public'
order by 4 desc;

--- 2.3.5 User List
SELECT username
  FROM pg_user;

```

2.4 Other Scripts and Artifacts

Copy any other scripts such as DML scripts etc. from your GIT instance and share them with the LeapLogic team to produce more extensive insights.

3. Getting Help

Contact LeapLogic technical support at info@leaplogic.io