

Migrate the PostgreSQL database created on Alibaba Cloud ECS to PostgreSQL of ApsaraDB for RDS

Alibaba Cloud Professional Service Team
July 2017

Contents

Introduction	3
Background information	3
Restrictions	3
Migration procedure	6
Further Reading	11

Introduction

To migrate your local database to ApsaraDB, you can use the data migration feature provided by Data Transmission Service (DTS) of Alibaba Cloud. The data migration feature of DTS helps you easily migrate your data to the cloud with just a few simple link configuration steps on the Data Transmission console.

This article introduces the solution of PostgreSQL database migration from Alibaba Cloud ECS to ApsaraDB for RDS., including the migration of data, views, user-defined functions, triggers and other objects.

This solution implements full migration of PostgreSQL database through DTS, from which you can learn how to migrate PostgreSQL database through DTS.

Background information

The migration types for PostgreSQL database supported by DTS are structure migration and full migration. The detailed information is as follows.

- Structure migration

Structure migration migrates the structure definition of the migrated object to the destination instance. For PostgreSQL database, the objects supported by DTS for structure migration include table, trigger, view, sequence, function, user-defined type, rule, domain, operation, and aggregate.

- Full migration

DTS migrates all the existing objects data of the migrated source database to the destination instance.

Restrictions

- Full migration supports the PostgreSQL versions of 9.2, 9.3, 9.4 and 9.5. Incremental migration supports the source PostgreSQL versions of 9.4.8 and 9.5.
- DDL operations are not supported during migration.
- C-language functions are not supported during migration.
- If the object name mapping function is enabled, other objects dependent on this object may fail to be migrated.

Migration sequence

To streamline the dependency between objects and improve migration success rates for PostgreSQL database migration, DTS defines the migration sequence of structure objects and data as follows:

1. Migrate structure objects: table, view, sequence, function, user-defined type, rule, domain, operation and aggregate.

2. Migrate all the existing data.

Note: After the full migration is complete, the migration progress in the task list is "Structure migration 100%" , "Full migration 100%" , and the migration status is "In migration". At this time, the migration task is migrating the structure objects in step 3. Do not end the task manually, otherwise it may cause loss of the migrated data.

3. Migrate structure objects: trigger and foreign key. [Environment preparation](#)

- For the destination database on RDS
 - a) Create a PostgreSQL instance on RDS.
 - b) Create and configure the database and account used for migration in the PostgreSQL instance of RDS.
 - c) Grant the account in the PostgreSQL instance of RDS with required permissions as follows:
 - ◆ For structure migration: the usage permission of pg_catalog.
 - ◆ For full migration: the select permission of the migrated objects.

- For the source database on ECS
 - a) Create an ECS instance with the operating system of Windows Server 2008 R2 and an Internet IP address.
 -
 - b) Install PostgreSQL 9.5 x64 on ECS.
 - c) Create the source PostgreSQL database, view, stored procedure, user-defined functions, and other resources.
 - d) Create an account used for migration and grant the account with required permissions as follows:
 - ◆ For structure migration: the create and usage permissions of the migrated objects.
 - ◆ For full migration: the owner permission of schema.

The command to create an account is as follows:

```
create user <username> password '<password>'
```

Parameter descriptions:

- ◆ username: The account to be created.
- ◆ password: The account logon password.

The command to grant an account is as follows:

```
GRANT <privileges> ON <tablename> TO <username>;
```

Parameter descriptions:

- ◆ privileges: The operation permission of the account, such as SELECT, INSERT and UPDATE. If you want to authorize all the permissions to the account, you can use ALL.
- ◆ tablename: The table name. If you want to authorize all the table permissions to the account, you can use wildcard character "*".
- ◆ username: The account name to be authorized.

Example:

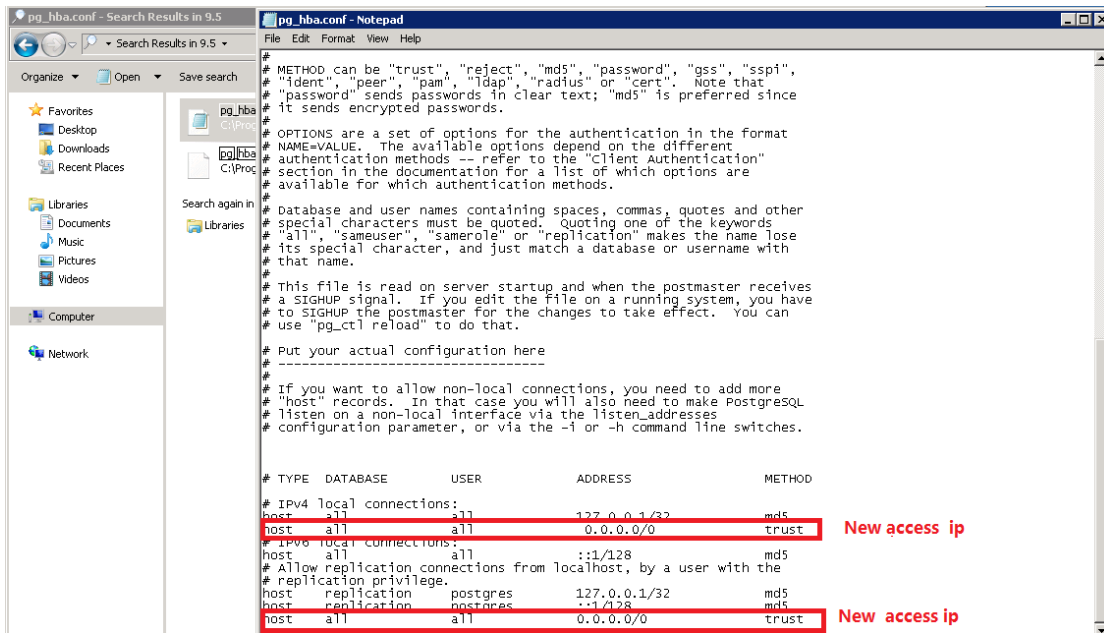
```
GRANT ALL ON * TO myuser;
```

- Modify the configuration file `pg_hba.conf`.

The client authentication is controlled by a configuration file. Its file name is usually `pg_hba.conf` and the file is stored in the data directory of the database cluster. Changes to this file only impact new connections established after the changes take effect, but have no effect on the current connection.

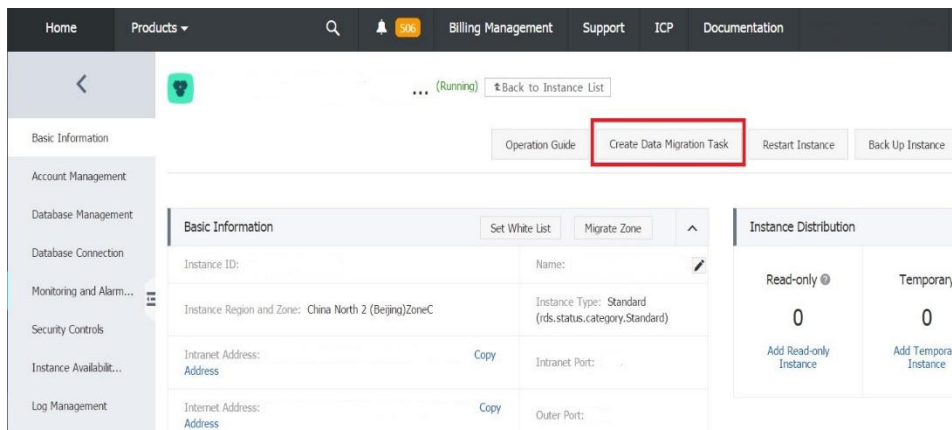
Create two new access records in the configuration file, as shown in the following figure. The contents marked with red boxes indicate the newly added IP access records. You must configure the authentication policies. Otherwise connection to the database will fail because the database in the authentication policies does not support remote access by default.

Alibaba Cloud - Migrate the PostgreSQL database created on Alibaba Cloud ECS to PostgreSQL of ApsaraDB for RDS



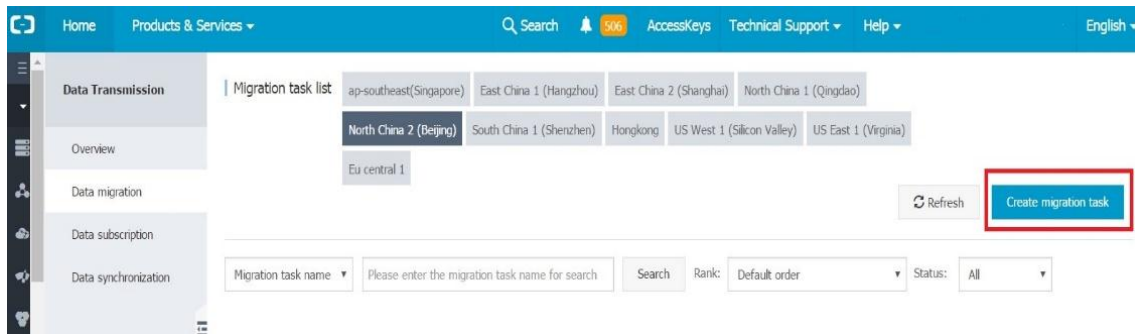
Migration procedure

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the instance ID to enter the **Basic Information** page.
4. Click Create **Data Migration Task** to log on to the Data Transmission console, as shown in the following figure.



5. In the left-side navigation pane, select **Data migration** to enter the **Migration task list** page.
6. Click **Create migration task** to enter the Data Transmission task configuration page, as shown in the following figure.

Alibaba Cloud - Migrate the PostgreSQL database created on Alibaba Cloud ECS to PostgreSQL of ApsaraDB for RDS



7. Enter the information of the source and target databases, as shown in the following figure.

The screenshot displays the 'Create migration task' form. At the top, there is a 'Task name' field. Below this, the form is divided into two main sections: 'Source database' and 'Target database'. The 'Source database' section includes fields for: '* Instance type:' (ECS-based databases), '* Instance region:' (ap-southeast(Singapore)), '* ECS instance ID:' (a blurred ID), '* Database engine:' (PostgreSQL), '* Port:' (3433), '* Database name:', '* Database account:', and '* Database password:'. A 'Test the connection' button is located at the bottom right of this section. The 'Target database' section includes fields for: '* Instance type:' (RDS instance), '* Instance region:' (ap-southeast(Singapore)), '* RDS instance ID:' (Select DRDS instance), '* Database name:', '* Database account:', and '* Database password:'. A 'Test the connection' button is also present at the bottom right of this section. Red boxes highlight the 'Task name' field and the configuration fields for both source and target databases.

Parameters description:

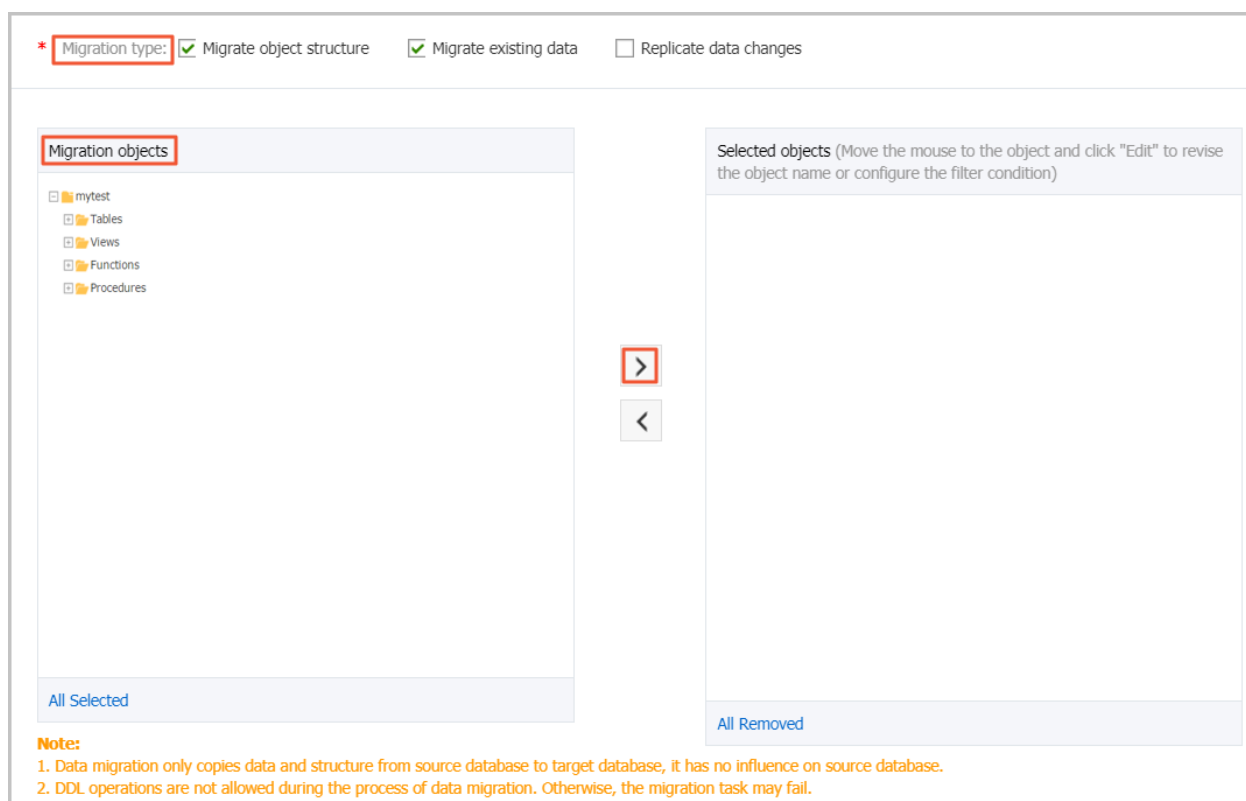
- Task name: By default, Data Transmission generates a name for every task automatically. The task name is not required to be unique. You can modify the name. A name indicating the specific services of the task is recommended to facilitate task identification.
- Source database
 - ◆ Instance type: Select **ECS-based databases**.
 - ◆ Instance region: Select the region where the ECS instance is located.
 - ◆ ECS instance ID: Select the ECS instance ID.

Alibaba Cloud - Migrate the PostgreSQL database created on Alibaba Cloud ECS to PostgreSQL of ApsaraDB for RDS

- ◆ Database engine: Select **PostgreSQL**.
 - ◆ Port: The listener port of the PostgreSQL instance.
 - ◆ Database account: The account used to access the PostgreSQL instance.
 - ◆ Database password: The password corresponding to the account used to access the PostgreSQL instance.
- Target database
- ◆ Instance type: Select **RDS instance**.
 - ◆ Instance region: Select the region where the RDS instance is located.
 - ◆ RDS instance ID: Select the ID of the target RDS instance to be migrated.
 - ◆ Database account: The account used to access the RDS instance.
 - ◆ Database password: The password corresponding to the account used to access the RDS instance.
8. Click **Test the connection** in the Source database and Target database areas respectively to test whether the connection information is correct.
9. If **Test passed** appears, click **Authorize whitelist and enter into next step** to enter **the Migration class and list** page.

Note: If the connection test fails, view the diagnosis and rectify the faults accordingly first.

10. Select the migration type and migration objects, as shown in the following figure.



Parameters description:

- Migration type
 - ◆ DTS supports structure migration and full migration for PostgreSQL database migration.
 - ◆ If you only need full migration, select **Migrate object structure** and **Migrate existing data**.
- Migration objects
 - ◆ The migration object can be a database, a table or a column.
 - ◆ By default, after the object is migrated to an RDS instance, the object name remains the same with that in the local PostgreSQL instance. If the object you migrate has different names on the source and destination instances, you need to use the object name mapping feature provided by DTS. Detailed usage can be found in [Database Table Column Mapping](#).

11. Click **Pre-check and start**.

Note: A pre-check will be performed before a migration task is formally started. Migration can be started only after the pre-check is successful. If the pre-check fails, check the failure details by clicking the button after the specific check items, rectify the faults accordingly, and perform a pre-check again.

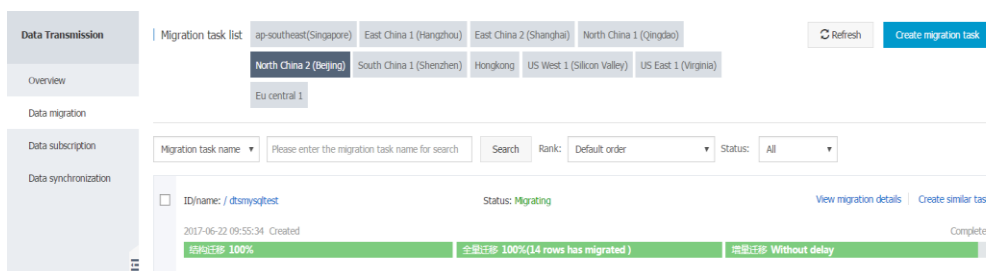
12. If the pre-check is successful, click **Buy and start now** to start the migration task.

13. Go back to the Data Transmission console.

14. In the left-side navigation pane, select **Data migration** to enter the **Migration task list** page.

15. Select the region where the target RDS instance is located.

16. Find the migration task and view the progress, as shown in the following figure. If you want to view the migration details, click **View migration details**.

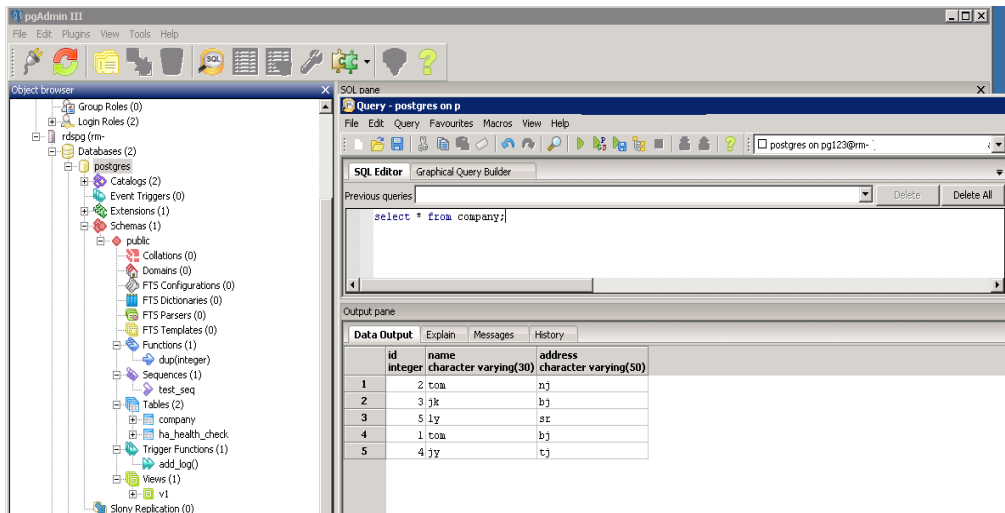


Data validation after migration

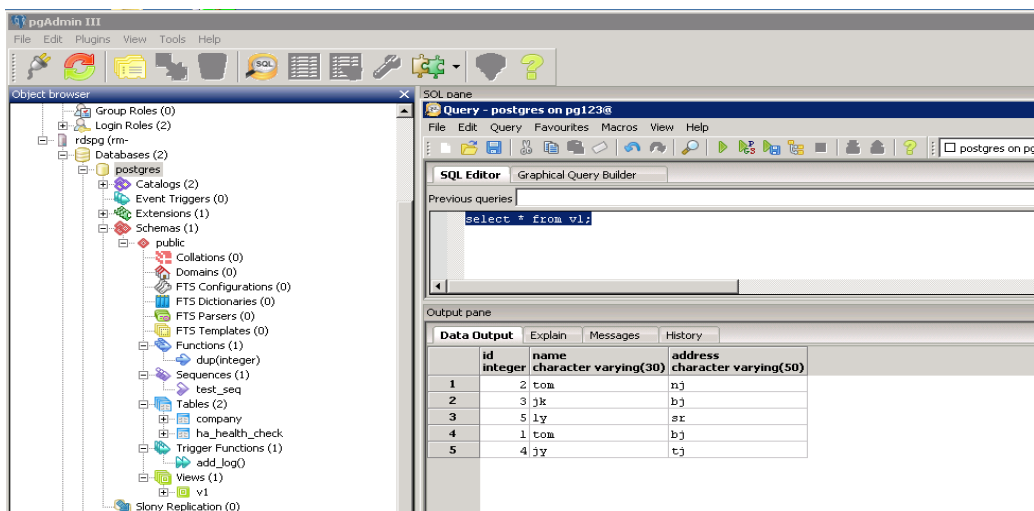
When a full migration task is completed, you can use the pgAdmin tool to connect to the destination database instance to check whether the data on it is consistent with that on the source database, and whether the migrated views, triggers, functions and other objects work normally.

- Data validation

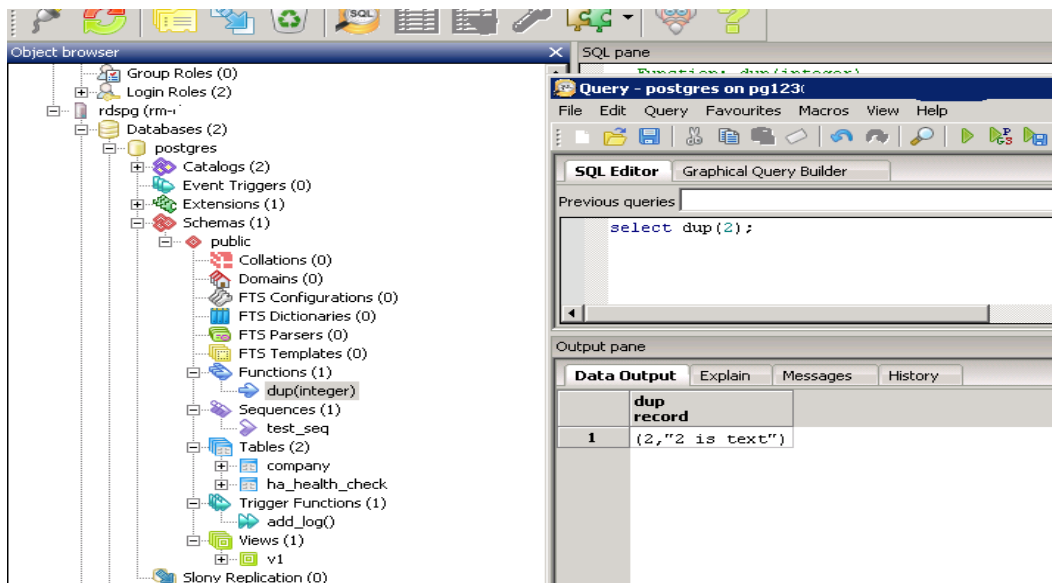
Alibaba Cloud - Migrate the PostgreSQL database created on Alibaba Cloud ECS to PostgreSQL of ApsaraDB for RDS



- Views

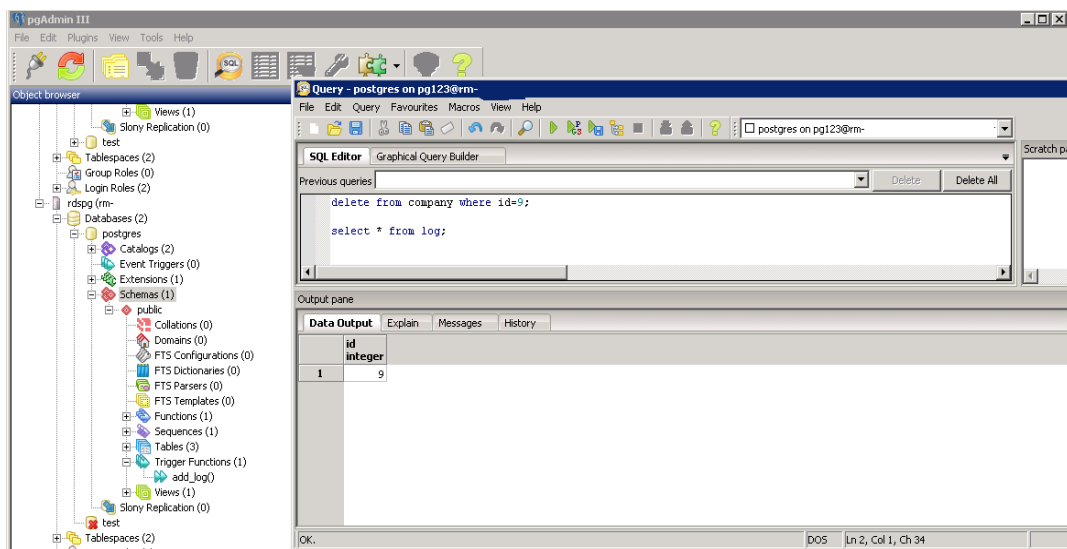


- User-defined functions



- Triggers

Alibaba Cloud - Migrate the PostgreSQL database created on Alibaba Cloud ECS to PostgreSQL of ApsaraDB for RDS



Further Reading

Migrate PostgreSQL database on Amazon RDS to PostgreSQL of ApsaraDB for RDS