

TencentDB for PostgreSQL

Getting Started

Product Documentation



Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Getting Started

- Creating TencentDB for PostgreSQL Instance

- Connecting to TencentDB for PostgreSQL Instance

- Managing TencentDB for PostgreSQL Instance

- Importing Data

- Migrating Data with DTS

Getting Started

Creating TencentDB for PostgreSQL Instance

Last updated : 2024-01-24 11:08:34

This document describes how to create a TencentDB for PostgreSQL instance in the console.

Prerequisites

You have [registered a Tencent Cloud account](#) and [completed identity verification](#).

Directions

1. Log in to the [TencentDB for PostgreSQL purchase page](#), specify the database instance information as needed, confirm that everything is correct, and click **Buy Now**.

Billing Mode: Pay-as-you-go billing is supported.

Region: The region where the instance is actually deployed. To minimize delay, we recommend the same region as the CVM instance to be connected to.

AZ: Physical IDCs where electric power facilities and networks are independent from each other within the same region. To minimize delay, we recommend the same AZ as the CVM instance to be connected to. Both multi-AZ deployment (the primary and standby nodes are in different AZs) and single-AZ deployment (the primary and standby nodes are in the same AZ) are supported. Specific primary and standby AZs are as displayed on the actual purchase page.

Network: The network where the instance is deployed. To minimize delay, we recommend the same network as the CVM instance to be connected to.

Note:

VPC: It is a logically isolated network space in Tencent Cloud. In a VPC, you can customize IP ranges, IP addresses, and routing policies.

Architecture: TencentDB for PostgreSQL supports the dual-server high-availability (one-primary-one-standby) architecture by default.

Database Version: The features available vary by PostgreSQL kernel version. For more information, see the official descriptions of PostgreSQL [10](#), [11](#), [12](#), and [13](#).

Instance Specification: Instance performance and base prices depend on its specification.

Disk: An SSD disk (local disk) is used by default.

Backup Space: You will receive 50% of the instance capacity for free to use as backup capacity. Any usage exceeding this complimentary capacity is currently free as well.

Instance Name: Name the instance now or later with up to 60 letters, digits, underscores, and hyphens.

Character Set: TencentDB for PostgreSQL supports UTF8 and LATIN1 character sets.

Username: The account name can contain 1–16 letters (case-insensitive), digits, or underscores. It cannot be `postgres` or start with a digit or `pg__`.

Password: The password can contain 8–32 characters. We recommend you use a password of at least 12 characters. It cannot start with a slash (/) and must contain all the following types of characters:

Lowercase letters (a–z)

Uppercase letters (A–Z)

Digits (0–9)

Special symbols `()~!@#$%^&*~+=_ | { } [] : ; ' < > , . ? /`

Project: If instances need to be managed by different teams, assign the instances to the projects of different teams accordingly.

Security Group: It serves as a stateful virtual firewall with filtering feature for configuring network access control for one or more TencentDB instances. It is an important network security isolation tool provided by Tencent Cloud.

Tag: It facilitates resource categorization and management.

Quantity: The number of instances that can be purchased at a time. To avoid faulty operations, an upper limit of 10 has been set for this parameter. If you want to purchase more instances, make multiple purchases.

Terms of Service: Read and click it. For more information, see [Terms of Service](#).

2. After the purchase is completed, you will be redirected to the [instance list](#). After the status of the instance changes to **Running**, the instance can be connected to.

Subsequent Operations

You can use a standard SQL client to connect to the TencentDB for PostgreSQL instance at its private or public network address. For more information, see [Connecting to TencentDB for PostgreSQL Instance](#).

Connecting to TencentDB for PostgreSQL Instance

Last updated : 2024-01-24 11:08:34

You can use a standard SQL client to access it over the private network or public network.

Private network access: a CVM instance can be used to access the private network address automatically assigned to the TencentDB instance. This access method relies on the high-speed private network of Tencent Cloud and features low delay. The two instances must be under the same account and in the same [VPC](#) in the same region or both must be in the classic network.

Note:

CVM and TencentDB instances in different VPCs (under the same or different accounts in the same or different regions) can be interconnected over private network through [Peering Connection](#).

Public network access: TencentDB for PostgreSQL can be accessed by using a public network address.

Note:

For public network access, the database instance's public IP needs to be enabled, which may expose your database service to attacks or intrusions on the public network. Therefore, it is recommended to log in to the database over the private network.

Public network access to TencentDB is suitable for development or auxiliary management of databases but not for business access in the production environment, as potentially uncontrollable factors may lead to unavailability of the public network access, such as DDoS attacks and bursts of high-traffic access.

Instances that currently support public network access are available only in Guangzhou, Shanghai, Beijing, Chengdu, Hong Kong (China), and Silicon Valley.

The following describes how to connect to a TencentDB for PostgreSQL instance from Windows and Linux CVM instances over the private and public networks.

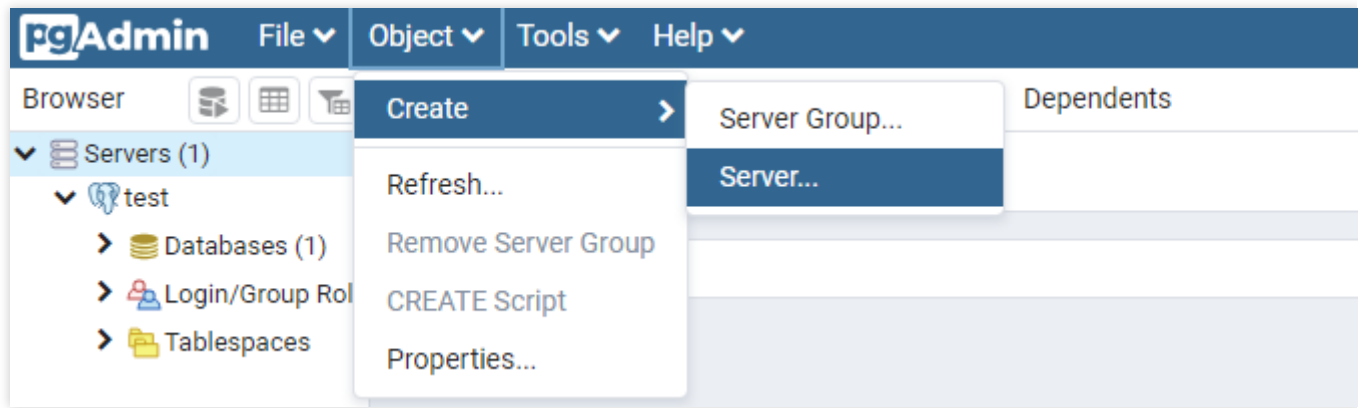
Connecting from a Windows CVM Instance

1. Download and install a standard SQL client in a Windows CVM instance or locally. For more information about how to log in to a CVM instance, please see [Customizing Windows CVM Configurations](#).

Note:

Take pgAdmin for example. You can download an installer [here](#) based on your operating system version.

2. In pgAdmin, select **Object > Create > Server**.



3. In the **Create - Server** dialog box, enter the information such as host name/address, port number, username, and password and then click **Save**.

Host name/address and port number: you can go to the [TencentDB for PostgreSQL console](#) and view them in **Private IPv4/IPv6 Address** or **Public IPv4/IPv6 Address** in the **Basic Info** section on the instance details page. If the public IP is not enabled, please see [Enabling Public Network Access](#).

Note:

Here, the private IP is VIP; database instances are accessed by connecting to the gateway cluster rather than the physical servers of database instances directly. Therefore, the private IP will remain unchanged in the event of server failures or primary/standby switchover.

Username and password: use the database admin username and password set when the instance is initialized. If you forget the password, you can go to the account management page in the [console](#) to reset it.

Create - Server

General Connection SSL SSH Tunnel Advanced

Host name/address

Port

5432

Maintenance database

postgres

Username

postgres

Password

Save password?

☐

Role

Service

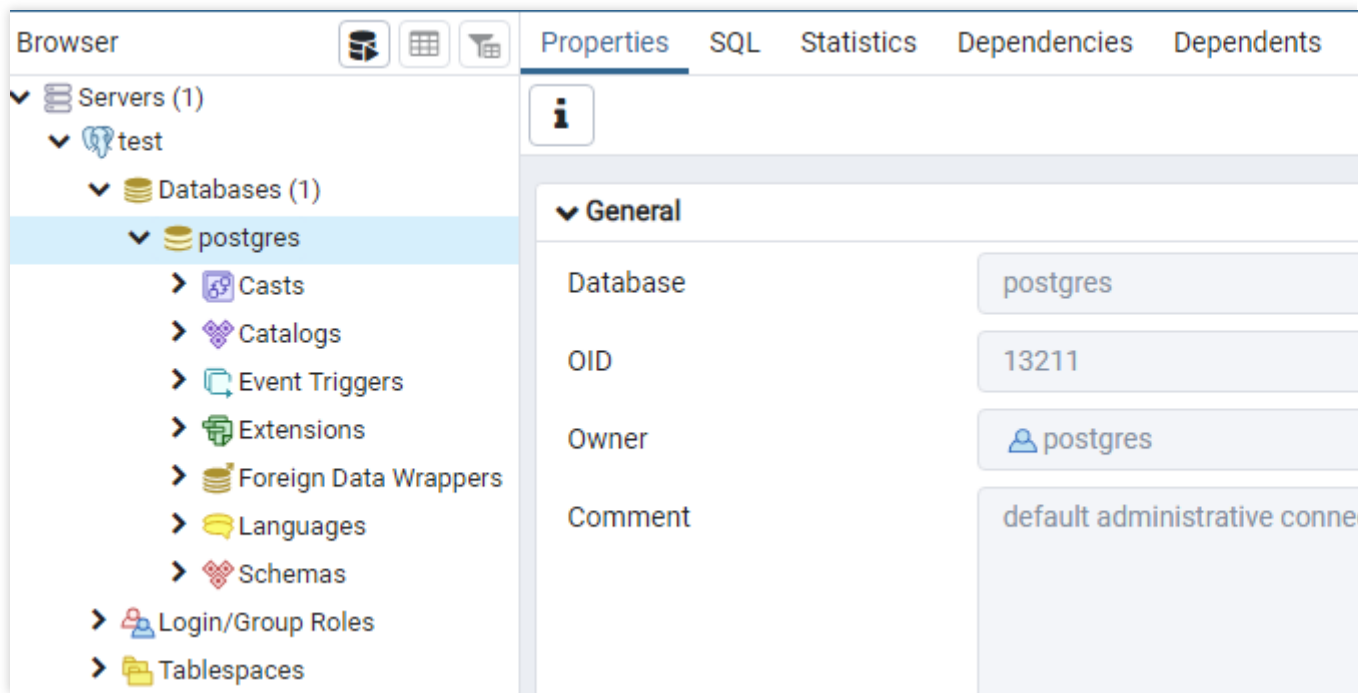
i ?

Cancel

Reset


Save

4. Then, select **Databases** > **postgres** on the left sidebar to view the connected server (database instance).



Connecting from a Linux CVM Instance

1. Install a psql client via yum in a Linux CVM instance or locally. For more information about how to log in to a CVM instance, please see [Customizing Linux CVM Configurations](#).
2. To install a psql client, please follow the instructions in [Restoring PostgreSQL Data on CVMs > Install PostgreSQL](#) to install PostgreSQL, as the psql client will be installed along with PostgreSQL.
3. Run the following command to log in to the PostgreSQL database:



```
psql -U username -h access address -p port -d postgres
```




Note:

To access from a CVM instance in the same VPC of the database instance, use the private IP of the database instance as the **access address**. To access from a Linux server on the internet, use the public IP of the database instance as the **access address**.

[Appendix. Enabling Public Network Access]

1. Log in to the [TencentDB for PostgreSQL console](#). In the instance list, click the instance ID/name or **Manage** in the **Operation** column to access the instance details page.
2. Click **Enable** next to **Public IPv4/IPv6 Address** in the **Basic Info** section on the instance details page.

Basic Info

Instance Name	Unnamed 
Instance ID	postgres-m7fgek8a 
Instance Status	Running
Region	Southwest China (Chengdu)
Availability Zone	Chengdu Zone 2
Network	Default-VPC - Default-Subnet
Project	Default Project Switch to another project
Character Set	UTF8
Private IPv4 Address	172.27.0.6:5432 
Public IPv4 Address	Enable
Tag	Modify

3. Click **OK** in the pop-up window and the request to enable public IP will be processed.
4. Once enabled successfully, the public IP can be found in the **Basic Info** section.


Managing TencentDB for PostgreSQL Instance

Last updated : 2024-01-24 11:08:34

Instance List Page

You can log in to the [TencentDB for PostgreSQL console](#) and enter the instance list page to view instance information and manage your instances.

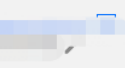

PostgreSQL - Instance List

 Chengdu (2)

Other regions (0) ▼

Create

More ▼

<input type="checkbox"/> Instance ID/Type/Name ▼	Monitoring / Status	Availability Zone	Configuration	Database Version	Billing Mode ⚙
<input type="checkbox"/> 	 Running	Chengdu Zone 2	Dual-Server High-Availability Edition 10 GB/2 GB Network: Default-VPC - Default-Subnet	PostgreSQL 12.4	Pay as you go

Instance restart

Note:

Exercise great caution when restarting a database, which plays a vital role in the business. Before the restart, we recommend you disconnect the database from server and stop writing data.

Restarting an instance does not change its physical attributes, so the public IP, private IP, and any data stored on the instance will remain unchanged.

After the restart, reconnection to the database is needed. Make sure your business has a reconnection mechanism.

Be sure to restart the instance during off-hours so as to ensure success and reduce impact on your business.

In the **Instance List**, find the instance to be restarted and click **More > Restart** in the **Operation** column. You can also select multiple instances and restart them in batches.

Note:

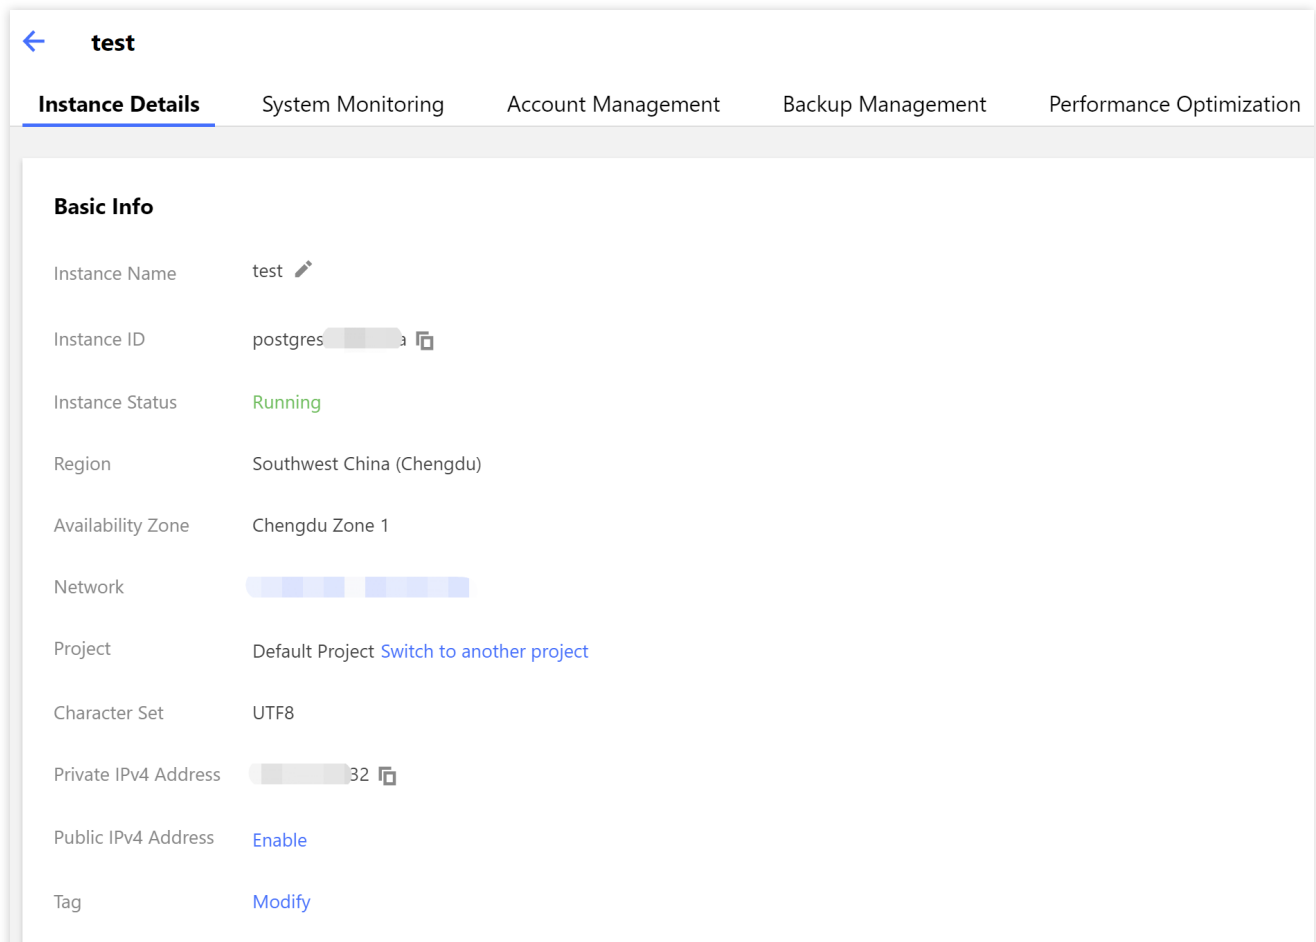
Generally, it takes tens of seconds to a few minutes to restart an instance, during which the instance cannot be accessed and existing connections to it will be closed.

Restart will fail if there are a large number of writes and dirty pages during the restart. In this case, the instance will roll back to the status before the restart and can still be accessed.

There is a chance of failure in restarting a database. If it takes more than 10 minutes to restart, you can [submit a ticket](#) for help.

Instance Management Page

After a TencentDB for PostgreSQL instance is created, click its ID in the [instance list](#) or click **Manage** in the **Operation** column to enter the instance management page, where you can view its details, monitor it, and manage its databases.



Instance details

On the **Instance Details** page, you can view and manage the basic information of the instance. The public network address is disabled by default; to use it, enable it manually.

System monitoring

On the **System Monitoring** page, you can view the monitoring data of various core metrics of the instance, including access, load, cache hit rate, SQL execution latency, XLOG sync delay, etc.

For more information on instance monitoring and alarming, see [Monitoring Feature](#) and [Alarming Feature](#).

Parameter settings

On the **Parameter Settings** page, you can modify parameters in batches or individually and query modification records. For more information, see [Setting Instance Parameters](#).

Account management

On the **Account Management** page, you can manage your account, such as modifying remarks and resetting the password.

Security group

On the **Security Group** page, you can query security group objects, configure security groups, and preview rules for instances. For more information, see [Managing Security Groups](#).

Note:

Security group outbound rules won't take effect for TencentDB for PostgreSQL, and inbound rules only apply to the private IP and port of TencentDB for PostgreSQL.

Backup management

On the **Backup Management** page, you can view and download backups and xlogs. For more information, see [Backing up Data](#).

Performance optimization

On the **Performance Optimization** page, you can view and download slow logs and error logs.

Read-Only instance

On the **Read-Only Instance** page, you can query and create read-only instances. You can also specify or create an RO group when creating a read-only instance.

Importing Data

Last updated : 2024-01-24 11:08:34

This document describes how to use DTS or PostgreSQL logical backups to restore data backup files to a TencentDB for PostgreSQL instance.

DTS

For detailed data migration solutions, see [Migrating Data with DTS](#). Currently, you can migrate many types of PostgreSQL databases to TencentDB for PostgreSQL.

Data Import/Export

Step 1. Prepare a TencentDB for PostgreSQL instance


[Purchase a TencentDB for PostgreSQL instance](#) and get its connection address in the [console](#).

Note:

Make sure that its character set is the same as that of the source instance.

Step 2. Make a logical backup of the source instance

1. Connect to the local (source) PostgreSQL database using a PostgreSQL client.
2. Run the following command to back up data:



```
pg_dump -U username -h hostname -p port -x databasename -f filename
```

The parameters are described as follows:

username: username of the local database.


hostname: hostname of the local database server. You can use "localhost" if you're logging in from the local database server.

port: port number of the local database.

databasename: name of the local database to be exported.

filename: name of the generated backup file.

-x: export data without object permission information of the source database. Importing data with permission information is prone to error. We recommend that you grant permissions later in the target database as needed. For example, if a database user named `pgtest` wants to back up a local PostgreSQL database, the user can log in to the PostgreSQL server and run the following command:




```
pg_dump -U pgtest -h localhost -p 4321 pg001 -f pg001.sql
```

Step 3. Restore data to the target instance

We recommend that you upload data to a CVM instance in a secure way (such as encryption and compression) and restore data to the target TencentDB for PostgreSQL instance over the private network.

1. Log in to the CVM instance.
2. On the PostgreSQL client, run the following command to import data to the target TencentDB for PostgreSQL instance:



```
psql -U username -h hostname -d databasename -p port -f filename
```

The parameters are described as follows:

username: database username of the target TencentDB for PostgreSQL instance


hostname: connection address of the target TencentDB for PostgreSQL instance

port: port number of the target TencentDB for PostgreSQL instance

databasename: database name of the target TencentDB for PostgreSQL instance

filename: name of the local backup file

For example:



```
psql -U pgtest -h 10.xxx.xxx.xxx -d pg001 -p 4321 -f pg001.sql
```

Because the permission configuration of the source database may be different from that of the target database, permission-related warnings or errors may appear during the data import process, which can be ignored.

Migrating Data with DTS

Last updated : 2024-01-24 11:08:34

In addition to PostgreSQL's `pg_dump` and `pg_restore` tools, you can also migrate data to the cloud or between database instances through Data Transmission Service (DTS), a migration tool provided by TencentDB.

Supported Features

Data can be migrated from self-built PostgreSQL to TencentDB for PostgreSQL.

Data can be migrated from PostgreSQL of other clouds to TencentDB for PostgreSQL.

Data can be migrated from CVM/container-based self-built PostgreSQL to TencentDB for PostgreSQL.

Data can be migrated from one instance to another in TencentDB for PostgreSQL.

Database can be migrated across accounts.

Data can be migrated between TDSQL-C for PostgreSQL and TencentDB for PostgreSQL.

TencentDB for PostgreSQL 10 and later supports cross-version migration, such as from PostgreSQL 10 to PostgreSQL 12 or PostgreSQL 14 to PostgreSQL 11.

PostgreSQL 9.4, 9.5, and 9.6 can be used as the source database to perform "full + incremental migration" by installing extensions on the source database; otherwise, only full migration is supported.

Notes

To avoid migration failure, we recommend that you read the notes below before migrating data.

When DTS performs full data migration, it will occupy certain source instance resources, which may increase the load of the source instance and the database pressure. If your database has low configurations, we recommend that you migrate data during off-peak hours.

When you migrate an instance over the public network, make sure that the source database service is accessible over the public network and keep the public network connection stable. If the network fluctuates or fails, migration will fail, and you need to restart the migration task.

During migration, the migration rate can be affected by factors such as the read performance of the source, the network bandwidth between the source and the target instances, and the specification of the target instance. Migration concurrency is determined by number of CPU cores of target instance, for instance, if the target instance has 2 cores, concurrency will be 2.

Correlated data objects need to be migrated together; otherwise, migration will fail. Common correlations include table reference by views, view reference by views, view/table reference by stored procedures/functions/triggers, and tables correlated through primary/foreign keys.

To ensure migration efficiency, the data of CVM-based self-built instances cannot be migrated across regions over the private network. If you need to migrate data across regions, you can do so over the public network .

To migrate the entire instance, there cannot be users and roles in the target database with the same name as those in the source database.

If you select “full + incremental migration”, tables in the source database must have a primary key; otherwise data inconsistency will occur in the source and target database. We recommend that you select full + incremental data migration for tables without the primary key.

The migration account of the source database needs to have the replication permission.

DDL sync is not supported during incremental migration. Meanwhile, large object modification is not recommended. To sync DDL, you need to execute DDL modification on both the source database and target database at the same time when data are same in both databases and execute DML after modification is completed.

If the source migration account is not `superuser` , select only the data to be migrated. We recommend that you manually create users and roles on the target database instead of migrating roles.

The `wal_level` parameter in the source database must be set to `logical` during incremental migration.

The value of `max_replication_slots` in the source database must be greater than the number of databases to be migrated during incremental migration.

The value of `max_wal_senders` in the source database must be greater than the number of databases to be migrated during incremental migration.

The value of `max_worker_processes` of the target database must be greater than that of `max_logical_replication_workers` during incremental migration .

Note:

If logical replication is performed on the source during the migration, or if slots or senders are occupied by tasks such as backup, the migration may fail; therefore, we recommend that you set larger values for the above parameters.

The available size of the target database space must be at least 1.2 times that of the instances to be migrated in the source database. Incremental data migration will execute UPDATE and DELETE operations, causing some tables in the database to generate data fragments. Therefore, after migration is completed, the size of the tables in the target database may be larger than that in the source database. This is mainly because that the `autovacuum` trigger conditions of the source and target databases are different. If the data in the source database is modified infrequently, the capacity of the target database may be smaller than that of the source instance.

The target database and the source database can't have migration object with the same name. To migrate roles, the account name can't be the same in both databases.

Incremental migration is not supported for special extensions such as timescaledb and pipelinedb.

To avoid migration failure by unsupported extensions, the extensions must be the same in both the source and the target.

Most extensions have no impact on the migration, but the pre-migration check will check whether there is a different extension in the target database. If yes, the extension will be created in the target database during migration. If it is not created successfully, the migration will fail.

The migration user in the source database must have permissions on all objects; otherwise the data export will fail. Ensure that the rules on character set and translation are same in the source and target database; otherwise the query result will be different. For details, see [Parameter Configuration Conflict Check](#). If language and collation is “C”, and the counterpart is “UTF8” in other database, you can ignore this note.

It is recommended that all tables in the source database have a primary key. If not, the migration efficiency will be reduced, and same data will occur in extreme scenarios.

Directions

1 (Optional) When PostgreSQL 9.4, 9.5, and 9.6 are used as the source database for “full + incremental migration”, you can install the `tencent_decoding` extension as instructed below. For other scenarios, skip this step.


1. Download extension based on architecture of server where source database resides.

Only “x86_64” and “aarch64” system architectures are supported.

The version of the extension need to be the same as that of PostgreSQL.


Requirements for the Glibc version: x86_64 system should be v2.17 - 323 or later; aarch64 system should be v2.17 - 260 or later.

[View Glibc version on Linux system](#)



```
RHEL/CentOS: rpm -q glibc
```

View Glibc version on other operation systems (Debian/Ubuntu/SUSE)



```
ldd --version | grep -i libc
```

Download address: [x86_64 9.4](#)、[x86_64 9.5](#)、[x86_64 9.6](#)、[aarch64 9.4](#)、[aarch64 9.5](#)、[aarch64 9.6](#)

2. Place the downloaded tencent_decoding.so file in the lib folder of the Postgres process directory without restarting the instance.

3. Log in to the [DTS console](#), select **Data Migration** on the left sidebar, click **Create Migration Task**, and enter the **Create Migration Task** page.

4. On the **Create Migration Task** page, select the types, regions, and specifications of the source and target instances and click **Buy Now**.

Configuration Item	Description
Source Instance Type	It is selected based on the source database type, and can't be modified after purchase. In this document, select PostgreSQL.
Source Instance Region	Select the source database region. If the source database is a self-built one, select a region nearest to it.
Target Instance Type	It is selected based on the target database type, and can't be modified after purchase. In this document, select PostgreSQL.
Target Instance Region	Select the target database region.
Specification	Select the specification of the migration linkage based on your business conditions.

5. Complete task configuration, source database settings, and target database settings on the “Set source and target databases” page. After the connectivity test for the source and target databases is passed, click **Create**.

Note:

If the connectivity test fails, troubleshoot as prompted or as instructed in [Database Connection Check](#) and try again.

1 Set source and target databases > 2 Set migration options and select migration objects > 3 Verify task

Task Configuration

Task Name *

Running Mode * Immediate execution Scheduled execution

Source Database Settings

Source Database Type * PostgreSQL

Access Type * Public Network Self-Build on CVM Direct Connect VPN Access Database CCN Intranet

Cross-/Intra-Account * Intra-account Cross-account

Region North China region (Beijing)

Database Instance * Please select

Account *

Password *

Test Connectivity

Target Database Settings

Target Database Type * PostgreSQL

Region North China (Beijing)

Configuration Type	Configuration Item	Description
Task Configuration	Task name	Set a task name that is easy to identify
	Running Mode	Immediate execution: The task will be started immediately after the task verification is passed. Scheduled execution: You need to configure a task execution time and the task will be started automatically then.
	Tag	The tag is used to manage resources by category from different dimensions. If the existing tag does not meet your requirements, you can manage tags in the console.
Source Database Settings	Source Database Type	The source database type selected during purchase, which can't be modified.

	Region	The source database region selected during purchase, which can't be modified.
	Access Type	<p>Select a type based on your scenario. In this document, TencentDB is selected as an example. For the preparations for different access types, see Overview.</p> <p>To ensure migration efficiency, the data of CVM-based self-built instances cannot be migrated across regions over the private network. If you need to migrate data across regions, you can do so over the public network.</p> <p>Public Network: The source database can be accessed through a public IP.</p> <p>Self-Built on CVM: The source database is deployed in a CVM instance.</p> <p>Direct Connect: The source database can be interconnected with VPCs through Direct Connect.</p> <p>VPN Access: The source database can be interconnected with VPCs through VPN Connection.</p> <p>Database: The source database is a TencentDB database.</p> <p>CCN: The source database can be interconnected with VPCs through CCN.</p>
	Database Instance	Select the instance ID of source PostgreSQL database .
	Account	The account of the source PostgreSQL database, which needs to have required permissions.
	Password	The password of the source PostgreSQL database account.

Target Database Settings	Target Database Type	The target database type selected during purchase, which cannot be changed.
	Region	The target database region selected during purchase, which cannot be changed.
	Access Type	In this document, “TencentDB” is selected by default based on your scenario.
	Database Instance	Select the target TencentDB instance ID.
	Account	The account of the target TencentDB database, which needs to have required permissions.
	Password	The password of the target database account.

6. On the “Set migration options and select migration objects” page, set the migration type and objects, and click **Save**.

Configuration Item	Description
Migration Type	<p>Select a type based on your scenario.</p> <p>Structural migration: Structured data such as databases and tables in the database will be migrated.</p> <p>Full migration: The entire database will be migrated. The migrated data will include only the existing data from the source database when the task is started, not incremental data written to the source database after the task is started.</p> <p>Full + incremental migration: The migrated data will include the existing data from the source database when the task is started as well as the incremental data written to the source database after the task is started. If there are data writes to the source database during migration, and you want to smoothly migrate the data in a non-stop manner, select this type.</p>
Migration Object	<p>Entire instance: Migrate the entire database instance, including the metadata definitions of roles and users but excluding system databases such as system objects in PostgreSQL.</p> <p>Specified objects: Migrate specified objects.</p>
Specified objects	Select the object to be migrated in the “Source Database Object” box, and drag it to the “Selected Object” box.

7. Verify the migration task on the “Verify task” page. After the task is verified, click **Start**.

If the verification failed, fix the problem as instructed in [Fix for Verification Failure](#) and initiate the verification task

again.

Failed: It indicates that a check item failed and the task is blocked. You need to fix the problem and run the verification task again.

Alarm: It indicates that a check item doesn't completely meet the requirements, and the task can be continued, but the business will be affected. You need to assess whether to ignore the alarm or fix the problem and continue the task based on the alarm message.

8. Return to the data migration task list, and the task will be in the “Preparing” status. After running for 1-2 minutes, the data migration task will be started.

Select **Structural migration** or **Full migration**: Once completed, the task will be stopped automatically.

Select **Full + Incremental migration**: After full migration is completed, the migration task will automatically enter the incremental data sync stage, which will not stop automatically. You need to click **Complete** to manually stop the incremental data sync. Then, the task will enter the **Completed** status. At this point, do not make any changes to the source and target databases. The backend will automatically align some objects with the source.

Manually complete incremental data sync and business switchover at appropriate time.

Check whether the migration task is in the incremental sync stage without any lag. If so, stop writing data to the source database for a few minutes.

Manually complete incremental sync when the data gap between the target and the source databases is 0 MB and the time lag between them is 0 second.

9. (Optional) If you want to view or delete a task, click the task and select the target operation in the **Operation** column. For more information, see [Viewing Task](#).

10. After the migration task status becomes **Task successful**, you can cut over the business. For more information, see [Cutover Description](#).