

Alibaba Cloud ApsaraDB for MySQL Quick Start for PostgreSQL

Issue: 20190716

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.








1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company, or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, special, or punitive damages, including lost profits arising from the use

or trust in this document, even if Alibaba Cloud has been notified of the possibility of such a loss.

5. By law, all the content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of the Alibaba Cloud website, product programs, or content shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates).
6. Please contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

| Style | Description | Example |
|---|--|--|
|  | This warning information indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results. |  Danger: Resetting will result in the loss of user configuration data. |
|  | This warning information indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results. |  Warning: Restarting will cause business interruption. About 10 minutes are required to restore business. |
|  | This indicates warning information, supplementary instructions, and other content that the user must understand. |  Notice: Take the necessary precautions to save exported data containing sensitive information. |
| | This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user. |  Note: You can use Ctrl + A to select all files. |
| > | Multi-level menu cascade. | Settings > Network > Set network type |
| Bold | It is used for buttons, menus, page names, and other UI elements. | Click OK. |
| <code>Courier font</code> | It is used for commands. | Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder. |
| <i>Italics</i> | It is used for parameters and variables. | <code>bae log list --instanceid Instance_ID</code> |
| [] or [a b] | It indicates that it is an optional value, and only one item can be selected. | <code>ipconfig [-all -t]</code> |

| Style | Description | Example |
|--|--|-------------------------------------|
| <code>{}</code> or <code>{a b}</code> | It indicates that it is a required value, and only one item can be selected. | <code>switch {stand slave}</code> |

Contents

| | |
|---|----|
| Legal disclaimer..... | I |
| Generic conventions..... | I |
| 1 Limits..... | 1 |
| 2 General procedure to use RDS..... | 2 |
| 3 Create an instance..... | 3 |
| 4 Initial configuration..... | 5 |
| 4.1 Set a whitelist..... | 5 |
| 4.2 Apply for an Internet address..... | 8 |
| 4.3 Create database and account..... | 11 |
| 5 Connect to an instance..... | 21 |
| 6 Read/write external data files using oss_fdw..... | 26 |

1 Limits

To guarantee instance stability and security, ApsaraDB for PostgreSQL has the following restrictions.

| Operations | RDS restrictions |
|------------------------------------|---|
| Modify database parameter settings | Currently it is not supported. |
| Database root permission | RDS does not offer the superuser permission. |
| Database backup | Data backup can only be performed through pg_dump. |
| Data migration | Data backed up through pg_dump can only be restored through psql. |
| Build database replication | <p>The system automatically builds the HA mode based on PostgreSQL stream replication.</p> <p>The PostgreSQL standby node is invisible and cannot be accessed directly.</p> |
| Restart the RDS instance | The instance must be restarted through the RDS console or Open APIs. |
| Network setting | If the access mode of the instance is safe connection mode, enabling net.ipv4.tcp_timestamps in SNAT mode is not allowed. |

2 General procedure to use RDS

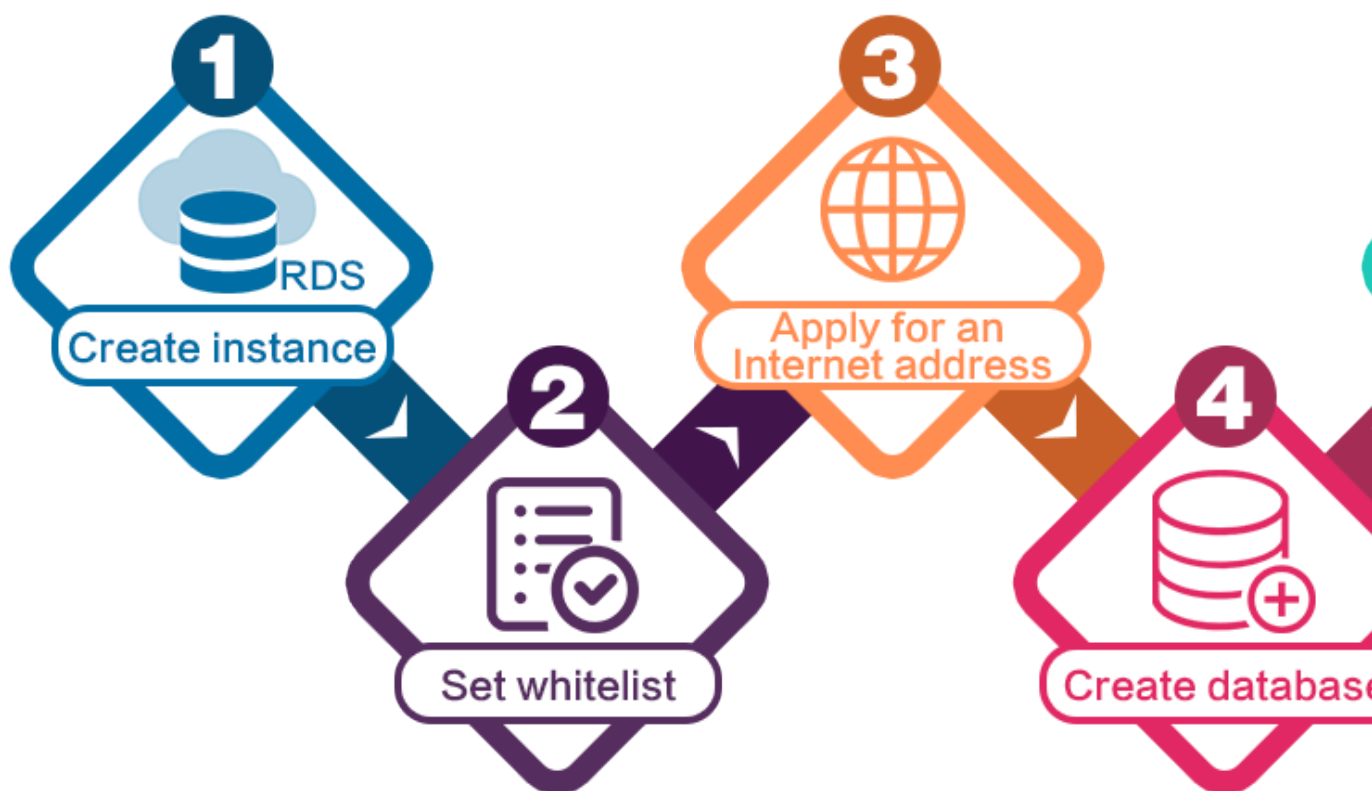
Purpose of the quick start

This document describes the procedure right from purchasing an RDS instance to using it. It also elaborates on how to create an ApsaraDB for RDS instance, perform basic settings, and connect to the instance database.

Quick start flowchart

If you use Alibaba Cloud ApsaraDB for RDS for the first time, see [Limits](#).

The following diagram explains the steps you must follow right from creating an instance to using it.



3 Create an instance

You can use the RDS console or APIs to create an RDS instance. For more information about instance pricing, see [Pricing of ApsaraDB for RDS](#). This document describes how to use the RDS console to create an instance. For more information about how to use APIs to create an instance, see [CreateDBInstance](#).

Prerequisites

- You must have registered to an Alibaba Cloud account.
- If you are creating a Pay-As-You-Go instance, make sure that your account balance is sufficient.

Procedure

1. Log on to the [RDS console](#).
2. On the Instances page, click Create Instance.
3. Select Subscription or Pay-As-You-Go. For more information about the billing method, see [Billing items and billing methods](#).
4. Select the instance configuration. The parameters are described as follows:
 - Basic configuration
 - Region and zone: Select the region and zone in which the instance is located. Some regions support single-zone and multi-zone instances, while some regions support only single-zone instances. For more information about regions and zones, see [Regions and zones](#).



Note:

Products in different regions cannot intercommunicate through the intranet, and you cannot change the instance region after creating an instance. Therefore, special attention is required when you select the region.

- Database engine: RDS supports MySQL, SQL Server, PostgreSQL, and PPAS . Different database types are supported in different regions. Choose the database type according to the instructions on the RDS console.
- Version: indicates the database version. Currently, RDS supports MySQL 5.5/5.6/5.7, SQL Server 2008 R2/2012, PostgreSQL 9.4, and PPAS 9.3. Different

database versions are supported in different regions. Choose the database version according to the instructions on the RDS console.

- Series: RDS instances support the Basic Edition, High-availability Edition, and Finance Edition. Different database versions support different series. Choose the instance series according to the instructions on the RDS console.
 - Network type: RDS supports the classic network and virtual private cloud (VPC). You can change the network type after creating an instance. For more information, see [Set network type](#).
 - Specifications: Specifications: indicate the CPU and memory occupied by the instance, the number of connections, and the maximum IOPS. For more information about instance specifications, see [Instance type list](#).
 - Storage: indicates space used by data, system files, binlog files, and transaction files.
 - Subscription time: indicates the duration of a Subscription instance.
 - Quantity: indicates the number of instances with the same configurations to be purchased.
5. Click Buy Now to enter the Confirm Order page.



Note:

To buy multiple instances with different configurations, click Add To List for each instance type and click Batch Purchase.

6. Select Product Terms of Service and Service Level Notice and Terms of Use, and then:
- Click Pay if the billing method of the instance is Subscription.
 - Click Activate if the billing method of the instance is Pay-As-You-Go.

4 Initial configuration

4.1 Set a whitelist

To ensure database security and stability, before using RDS instances, you must whitelist the IP addresses or IP address segments that need to access the database. We recommend that you periodically check and adjust your whitelist according to your requirements to maintain RDS security. This document provides information about and the procedure of setting a whitelist.

Background

You can access the RDS instance through the intranet, the Internet, or both the intranet and Internet. For more information about the applicable scenarios of each connection type (intranet and Internet), see Background of [Set intranet and Internet addresses](#).

- Access the RDS instances through the intranet.
- Access the RDS instances through the Internet.
- Access the RDS instances through both the intranet and Internet.

Before setting the connection type, you must add the IP addresses or IP address segments of your application service or the ECS instance to the whitelist of your RDS instance. When the whitelist is set, the system automatically generates the intranet IP address for the RDS instance. If you need an Internet IP address, refer to [Apply for an Internet address](#).



Note:

If you cannot connect to the RDS instance after adding the application service IP address to the whitelist, you must obtain the actual IP address of the application service.

Attention

- The system automatically creates a default whitelist group for each newly created RDS instance. This default whitelist group can only be modified or cleared, but cannot be deleted.

- For each newly created RDS instance, the local loopback IP address 127.0.0.1 is added to the default whitelist group by default. This means that all the IP addresses or IP address segments are prohibited to access this RDS instance. Therefore, you must delete 127.0.0.1 from the default whitelist group before you add other IP addresses or IP address segments to the RDS whitelist.
- % or 0.0.0.0/0 indicates any IP address is allowed to access the RDS instance. This configuration greatly reduces the security of the database and is not recommended.
-


Procedure

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the name of the target instance to go to the Basic Information page.
4. Select Security Controls in the left-side navigation pane to visit the Security Controls page.
5. On the Whitelist Settings tab page, click Modify of the default whitelist group, as shown in the following figure.



Note:

If you want to add a customized whitelist group to the RDS instance, you can click Clear of the default whitelist group to delete the IP address 127.0.0.1 first, and then click Add a Whitelist Group. The setting steps for a customized whitelist are similar to the following steps.



6. On the Modify Group page, add the IP addresses or IP address segments to access the RDS instance to the Whitelist field. If you want to add the ECS intranet IP addresses, click Upload ECS Intranet IP Address and select the IP addresses according to the prompt window, as shown in the following figure.



Note:

After you add a new IP address or IP address segment to the default group, the loopback address 127.0.0.1 is automatically deleted.

Group Name: default

Whitelist: 127.0.0.1

[Upload ECS Intranet IP Address](#) You can add 999 whitelists more

Specified IP address: Add an IP address to allow this IP to access RDS.
Specified IP segment: Add an IP segment to allow all the IP addresses in this segment to access RDS.
When you add multiple IP addresses, separate them by a comma (no space after the comma), such as "192.168.0.1,192.168.0.1/24".
[How to locate the local IP address](#)

White list will be effect after 1 minute

Parameters description:

- Group Name : It contains 2 to 32 characters including lowercase letters, digits, or underscores (_). The group name must start with a lowercase letter

and end with a letter or digit. This name cannot be modified once the whitelist group is successfully created.

- **Whitelist** : Enter the customized IP addresses or IP segments that are allowed to access the RDS instance.
 - If you enter an IP address segment, such as 10.10.10.0/24, it indicates that any IP address in the format of 10.10.10.X can access the RDS instance.
 - If you want to enter multiple IP addresses or IP address segments, separate them by commas (,) (do not add blank spaces), such as 192.168.0.1,172.16.213.9.
 - For each whitelist group, up to 1,000 IP addresses or IP address segments can be set for MySQL, PostgreSQL, and PPAS instances and up to 800 can be set for SQL Server instances.
- **Upload ECS intranet IP Address** : By clicking this button, you can select the intranet IP address of the ECS instance under the same account as the RDS instance. This is a quick method to add ECS intranet IP address.

7. Click OK.

Modify or delete the whitelist group


You can modify or delete the whitelist group according your business requirements. The detailed procedure is as follows:

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the name of the target instance to go to the Basic Information page.
4. Select Security in the left-side navigation pane.
5. On the Whitelist Settings tab page, click Modify or Delete of the target whitelist group.
6. Click OK after you modify the IP addresses or IP address segments. Alternatively, click Confirm if you are sure that the whitelist group is to be deleted.

4.2 Apply for an Internet address

RDS provides two types of addresses: intranet addresses and Internet addresses.

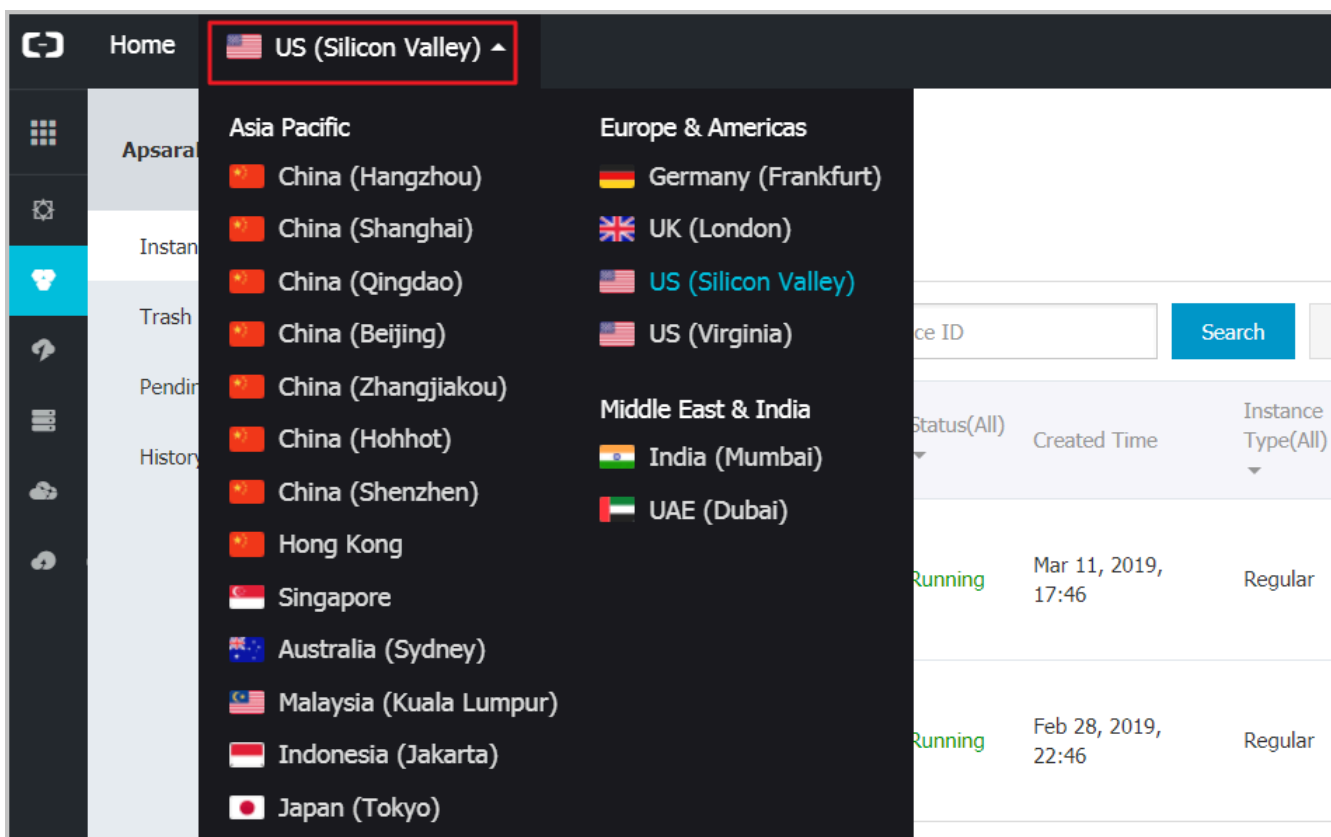
Intranet and Internet addresses

| Address Type | Description |
|------------------|--|
| Intranet address | <p>An intranet address is generated by default.</p> <p>Use the intranet address if all of the following conditions are met:</p> <ul style="list-style-type: none">• Your application is deployed on an ECS instance.• The ECS instance is located in the same region as your RDS instance.• The ECS instance has the same network type as your RDS instance. <p>We recommend that you use the intranet address to access your RDS instance because this is more secure and delivers optimal performance.</p> |
| Internet address | <p>You need to manually apply for an Internet address. You can also release it anytime.</p> <p>Use the Internet address if you cannot access RDS through the intranet. Specific scenarios are as follows:</p> <ul style="list-style-type: none">• An ECS instance accesses your RDS instance but the ECS instance is located in a different region or has a network type different from your RDS instance.• A server or computer outside Alibaba Cloud accesses your RDS instance. <div> Note:<ul style="list-style-type: none">• The Internet address and traffic are currently free of charge.• Using the Internet address reduces security. Please exercise caution• To ensure high security and performance, we recommend that you migrate your application to an ECS instance that is in the same region and has the same network type as your RDS instance and then use the intranet address.</div> |

Apply for an Internet address

1. Log on to the [RDS console](#).


2. In the upper-left corner, select the region where the RDS instance is located.



3. Find the RDS instance and click its ID.

4. In the left-side navigation pane, choose Connection Options.

5. Click Apply for Internet Address.



6. In the displayed dialog box, click OK.

The Internet address is generated.



Note:

You can view the Internet address only after the [whitelist](#) is configured.

7. (Optional) To modify the Internet address or port number, click **Modify Connection Address**. In the displayed dialog box, set the Internet address and port number and click **OK**.

- **Connection Type:** Select Internet address.



Note:

This option is available only after you have applied for the Internet address.

- **Connection Address:** You can modify the address prefix, which consists of 8 to 30 characters, including letters and digits, and starts with a lower-case letter.
- **Port:** The port number can be modified only if the RDS network type is classic network.

APIs

| API | Description |
|--|--|
| AllocateInstancePublicConnection | Applies for an Internet address for an instance. |

4.3 Create database and account

Before using RDS, you must create databases and accounts for the RDS instance. For PostgreSQL instances, you must create an initial account on the RDS console, and then you can create and manage the databases through a client. This document takes the pgAdmin 4 client as an example to introduce how to create databases and accounts for PostgreSQL instances.

Background information

- Databases under a single instance share all the resources of this instance. Each PostgreSQL instance supports one initial account, countless general accounts, and countless databases. You must create and manage the general accounts and databases through SQL statements.
- To migrate your local database to the RDS instance, you must create the same databases and accounts for the RDS instance as your local database.
- When assigning account permissions for each database, follow the minimum permission' principle and consider service roles to create accounts. Alternatively, rationally assign read-only and read/write permissions. When necessary, you can split accounts and databases into smaller units so that each account can only access data for its own services. If the account does not need to write data to a database, assign the read-only permission for the account.
- For database security, set strong passwords for the accounts and change the passwords regularly.

Procedure

1. Log on to the [RDS console](#).
2. Select the region where the target instance is located.
3. Click the ID of the instance to visit the Basic Information page.
4. In the left-side navigation pane, select Accounts.
5. Click Create Initial Account.

6. To create an account, set the related fields.

Create Account [Back to Account Management](#)

Database Account:

Your account name can have 2 to 16 characters including lower-case letters, digits, or underscores. It must begin with a letter and end with a letter or a digit.

***Password:**

Your password can have 8 to 32 characters including at least three of the following:

- Capital letters
- Lower-case letters
- Digits
- Special characters (!@#\$%^&*()_-=)

***Re-enter Password:**

Up to 1 accounts can be created.

Parameters description:

- **Database Account** : refers to the name of the initial account. It contains 2 to 16 characters including the lower-case letters, digits, or underscores (_). It must begin with a letter and end with a letter or digit.
- **Password** : refers to the password of the initial account. It contains 8 to 32 characters including at least three of the following: capital letters , lower-case letters, digits, and special characters (!@#\$%^&*()_-=)
- **Re - enter Password** : Re-enter the password to make sure the password is entered correctly.


7. Click OK.

8. Add the IP address that is allowed to access the RDS instance to the RDS whitelist.


For more information about how to set the whitelist, see [Set the whitelist](#).

9. Start the pgAdmin 4 client.

10. Right-click Servers, and then select Create > Server, as shown in the following figure.



11. On the General tab of Create - Server window, enter server name, as shown in the following figure.



The screenshot shows a window titled "Create - Server" with a close button in the top right corner. Below the title bar are two tabs: "General" (selected) and "Connection". The "General" tab contains the following fields:

- Name:** A text input field with a red rectangular highlight around its label.
- Server group:** A dropdown menu showing "Servers" with a downward arrow.
- Connect now?:** A checkbox that is checked.
- Comments:** A large, empty text area.

At the bottom of the window is a toolbar with four buttons: an information icon (i), a help icon (?), a "Save" button with a floppy disk icon, a "Cancel" button with a red X icon, and a "Reset" button with a circular arrow icon.

12. Click the Connection tab, and enter the information of the instance to be connected, as shown in the following figure.

The screenshot shows a 'Create - Server' dialog box with a 'Connection' tab selected. The dialog has a blue header bar with a close button. Below the header, there are two tabs: 'General' and 'Connection'. The 'Connection' tab is active, showing several input fields and a checkbox. The fields are: 'Host name/address', 'Port', 'Maintenance database' (with 'postgres' entered), 'Username', 'Password', 'Save password?' (with an unchecked checkbox), 'Role', and 'SSL mode' (with 'Prefer' selected in a dropdown). A red error message at the bottom states: 'Port' must be greater than or equal to 1024. At the bottom of the dialog are three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (yellow).

| Field | Value |
|----------------------|--------------------------|
| Host name/address | |
| Port | |
| Maintenance database | postgres |
| Username | |
| Password | |
| Save password? | <input type="checkbox"/> |
| Role | |
| SSL mode | Prefer |

'Port' must be greater than or equal to 1024.

Buttons: Save, Cancel, Reset

Parameters description:

- Host name / address : refers to the connection address of the RDS instance. If your application accesses the RDS instance through the intranet, enter the intranet IP address of the RDS instance. If your application accesses

the RDS instance through the Internet, enter the Internet IP address of the RDS instance. You can view the connection address and port number as follows:

- a. Log on to the [RDS console](#).
 - b. Select the region where the target instance is located.
 - c. Click the ID of the instance to visit the Basic Information page.
 - d. View the intranet and Internet IP addresses and ports in the Basic Information area.
- Port : refers to the port number of the the RDS instance. If your application accesses the RDS instance through the intranet, enter the intranet port number of the RDS instance. If your application accesses the RDS instance through the Internet, enter the Internet port number of the RDS instance.
 - Username : refers to the initial account name of the RDS instance.
 - Password : refers to the password of the initial account of the RDS instance.

13. Click Save.

14. If the connection information is correct, select Servers > server name > Databases > postgres. The following interface is displayed, which indicates that the connection to RDS instance is successful.




Note:


postgres is the default system database of the RDS instance. Do not do any operation in this database.



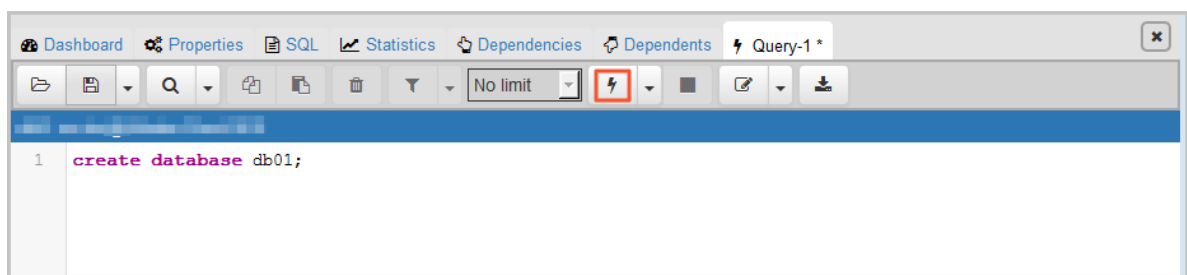
15. Click postgres, and then select Tools > Query Tool, as shown in the following figure.



16. Enter the following command on the Query-1 tab page to create a database, as shown in the following figure.




17. Click Execute/Refresh, as shown in the following figure.




If the execution is successful, the new database is created successfully.

18.Right-click Databases and click Refresh, and then you can find the newly created database, as shown in the following figure.

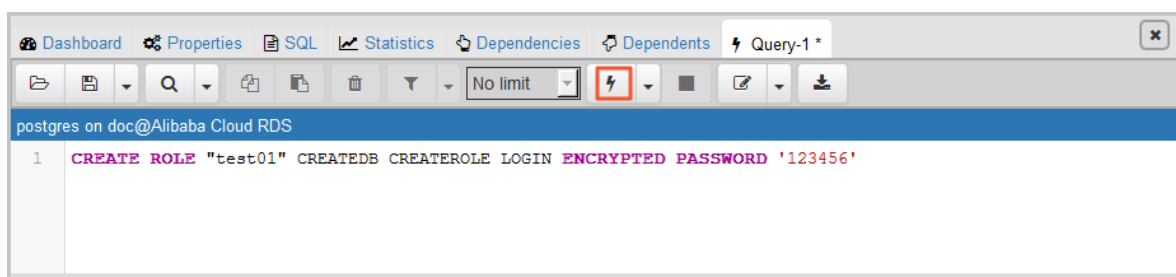


19.Enter the following command on the Query-1 tab page to create an account, as shown in the following figure.

```
CREATE ROLE "username" CREATEDB CREATEROLE LOGIN
ENCRYPTED PASSWORD 'password';
```

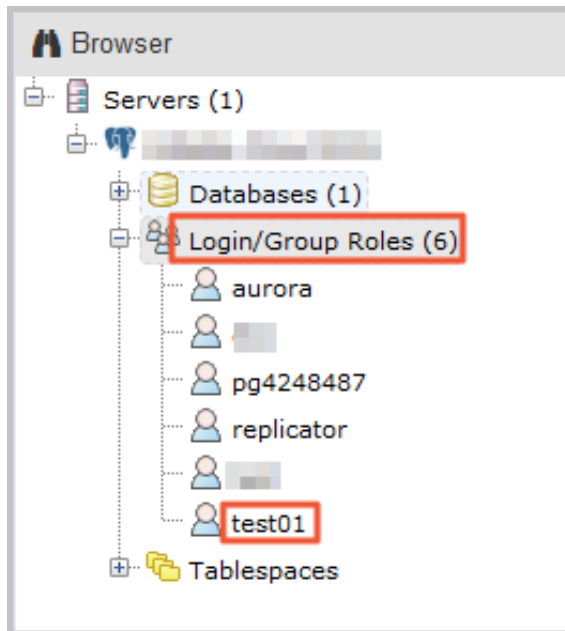


20.Click Execute/Refresh, as shown in the following figure.



If the execution is successful, the new account is created successfully.

21. Right-click Login/Group Roles and click Refresh, and then you can find the newly created account, as shown in the following figure.



5 Connect to an instance

You can connect to an RDS instance through the PostgreSQL client. This document introduces the connection procedure by taking the pgAdmin 4 client as an example.

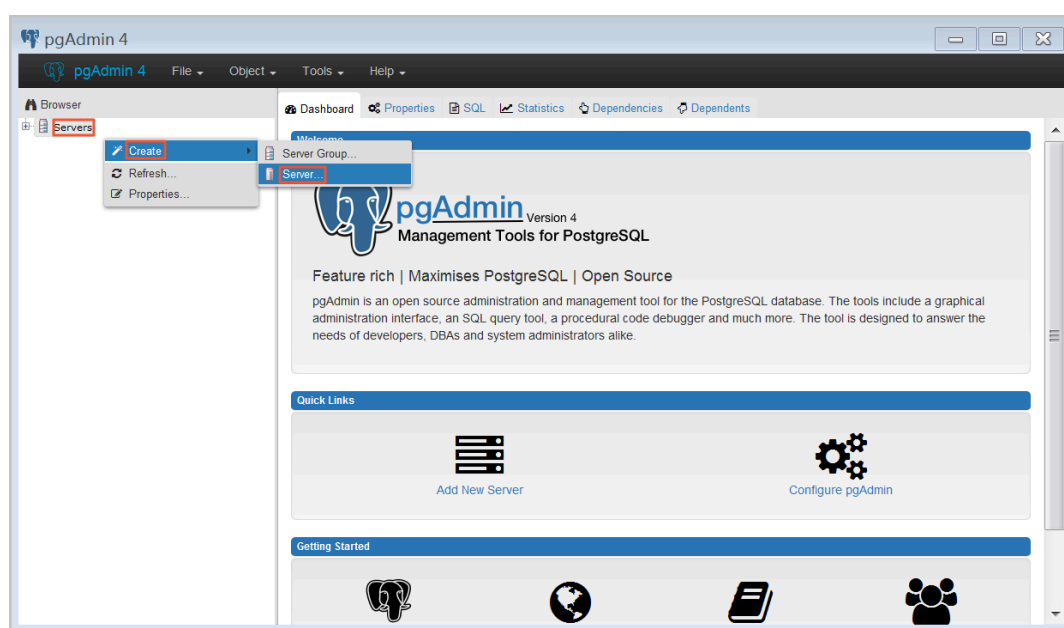
Background information

RDS for PostgreSQL is fully compatible with PostgreSQL, so you can connect to RDS in the way you connect to an on-premises PostgreSQL database. This document takes the pgAdmin 4 client as an example to introduce how to connect to an RDS instance. You can also adopt this method when using other clients. When you connect to an RDS instance through a client, choose to use an [intranet or Internet address](#) as follows:

- Use the intranet IP address when your client is installed on the ECS that is located in the same region and the same network type as the RDS instance to be connected.
- Use the Internet IP address for the other situations.

Log on through a client

1. Add the IP address that is allowed to access the RDS instance to the RDS whitelist.
For more information, see [Set the whitelist](#).
2. Start the pgAdmin 4 client.
3. Right click Servers, and then select Create > Server, as shown in the following figure.



4. On the General tab of Create - Server window, enter server name, as shown in the following figure.

The screenshot shows a 'Create - Server' dialog box with two tabs: 'General' and 'Connection'. The 'General' tab is active. It contains the following fields and controls:

- Name:** A text input field, highlighted with a red rectangular box.
- Server group:** A dropdown menu currently showing 'Servers'.
- Connect now?:** A checkbox that is checked.
- Comments:** A large text area for additional notes.

At the bottom of the dialog, there is a bar with an information icon, a help icon, and three buttons: 'Save' (blue), 'Cancel' (red), and 'Reset' (yellow).

5. Click the Connection tab, enter the information of the instance to be connected, as shown in the following figure.

Create - Server

General Connection

Host name/address

Port

Maintenance database: postgres

Username

Password

Save password? ☐

Role

SSL mode: Prefer

'Port' must be greater than or equal to 1024.

Save Cancel Reset

Parameters description:

- Host name / address : refers to the connection address of the RDS instance. If your application accesses the RDS instance through the intranet, enter the intranet IP address of the RDS instance. If your application accesses the RDS instance through the Internet, enter the Internet IP address of the RDS

instance. Perform the following steps to find the connection address and port number of the RDS instance.

- a. Log on to the [RDS console](#).
- b. Select the region where the target instance is located.
- c. Click the ID of the instance to visit the Basic Information page.
- d. In the Basic Information area, you can find the connection addresses and port numbers of the RDS instance.

- Port : refers to the port number of the RDS instance. If your application accesses the RDS instance through the intranet, enter the intranet port number of the RDS instance. If your application accesses the RDS instance through the Internet, enter the Internet port number of the RDS instance.
- Username : refers to the initial account name of the RDS instance.
- Password : refers to the password of the initial account name of the RDS instance.

6. Click Save.

7. If the connection information is correct, select Servers > server name > Databases > postgres. The following interface is displayed, which indicates that the connection to RDS instance is successful.



Note:

Postgres is the default system database of the RDS instance. Do not perform any operation in this database.



6 Read/write external data files using oss_fdw

In Alibaba Cloud, you can use oss_fdw plugin to load data on OSS to PostgreSQL and PPAS databases, and you can also write data in a database to OSS.

oss_fdw parameters

Similar to other fdw interfaces, oss_fdw can encapsulate data stored on OSS (external data sources), allowing you to read files on OSS. The process is like reading data from a table. oss_fdw provides unique parameters used for connecting to and parsing file data on OSS.



Note:

- Currently, oss_fdw can read and write the following file types in OSS: text/csv files and text/csv files in GZIP format.
- The value of each parameter needs to be quoted and cannot contain any useless spaces.

CREATE SERVER parameters

- **ossendpoint**: Address (host) used to access OSS from the intranet
- **id**: OSS account ID
- **key**: OSS account key
- **bucket**: OSS bucket, assigned after an OSS account is created

The following parameters are related to error tolerance in import and export modes. If network connectivity is poor, you can adjust these parameters to facilitate successful imports and exports.

- **oss_connect_timeout**: Connection expiration time, measured in seconds. Default value: 10s.
- **oss_dns_cache_timeout**: DNS expiration time, measured in seconds. Default value: 60s.
- **oss_speed_limit**: Minimum tolerable rate. Default value: 1,024 byte/s (1 Kbit/s).
- **oss_speed_time**: Maximum tolerable time. Default value: 15s.

If the default parameter values are used, a timeout error occurs when the transmission rate is smaller than 1 Kbit/s for 15 consecutive seconds.

CREATE FOREIGN TABLE parameters

- **filepath:** File name including a path on OSS.
 - A file name contains a path but not a bucket name.
 - This parameter matches multiple files in the corresponding path on OSS, and supports file loading to a database.
 - Files named in the format of filepath or filepath.x can be imported to a database. x in filepath.x must start from 1 and be consecutive, for example, filepath, filepath.1, filepath.2, filepath.3, and filepath.5.

The first four files are matched and imported, but the file named filepath.5 is not
- **dir:** Virtual directory on OSS.
 - dir must end with a slash (/).
 - All files (excluding subfolders and files in subfolders) in the virtual directory indicated by dir are matched and imported to a database.
- **prefix:** Prefix of the path in the data file. Regular expressions are not supported. You can set only one of the these parameters: prefix, filepath, and dir.
- **format:** File format, which can only be CSV currently.
- **encoding:** File data encoding format. It supports common PostgreSQL encoding formats, such as UTF-8.
- **parse_errors:** Parsing in error tolerance mode. The errors that occur during the file parsing process are ignored by row.
- **delimiter:** Delimiter specified for columns.
- **quote:** Quote character for a specified file.
- **escape:** Escape character for a specified file.
- **null:** Used to nullify the column matching a specified string. For example, null 'test' is used to set the column whose value is 'test' to null.
- **force_not_null:** Used to un-nullify the value of one or more columns. For example, force_not_null 'id' is used to set the values of the 'id' column to empty strings.
- **compressiontype:** Used to set whether the file read or written on OSS is compressed and set the compression format. Value range:
 - none: Uncompressed (default value)
 - gzip: compressed gzip file

- **compressionlevel**: Used to set the compression level of the compression format written to OSS, ranging from 1 to 9. The default value is 6.

**Note:**

- **filepath** and **dir** need to be specified in the **OPTIONS** parameter.
- Either **filepath** and **dir** must be specified, and they cannot be specified at the same time.
- The export mode currently only supports virtual folders, that is, only **dir** is supported.

Export mode parameters for CREATE FOREIGN TABLE

oss_flush_block_size and **oss_file_max_size** are added for the export mode.

- **oss_flush_block_size**: Buffer size for the data written to OSS at a time. Its default value is 32 MB, and the value range is 1 MB to 128 MB.
- **oss_file_max_size**: Maximum file size for the data written to OSS (subsequent data is written in another file when the maximum file size is exceeded). Its default value is 1,024 MB, and the value range is 8 MB to 4,000 MB.
- **num_parallel_worker**: The number of parallel compression threads in the compression mode in which the OSS data is written, ranging from 1 to 8. Its default value is 3.

**Note:**

oss_flush_block_size and **oss_file_max_size** are invalid for the import mode.

Auxiliary function

FUNCTION oss_fdw_list_file (relname text, schema text DEFAULT 'public')

- Used to obtain the name and size of the OSS file that an external table matches.
- The unit of file size is byte.

```
select * from oss_fdw_list_file (' t_oss ');
      name | size
-----
oss_test / test . gz . 1 | 739698350
oss_test / test . gz . 2 | 739413041
oss_test / test . gz . 3 | 739562048
```

```
( 3 rows )
```

Auxiliary feature

oss_fdw.rds_read_one_file: In read mode, it is used to specify a file that matches the external table. Once it is set, the external table matches only one file that is set during data import.

For example, set `oss_fdw.rds_read_one_file = 'oss_test/example16.csv.1'` ;

```
set oss_fdw . rds_read_one_file = ' oss_test / test . gz . 2 ' ;
select * from oss_fdw_list_file (' t_oss ');
      name | size
-----
 oss_test / test . gz . 2 | 739413041
( 1 rows )
```

oss_fdw example

```
# ( PostgreSQL ) Create the plugin
create extension oss_fdw ; ---- For PPAS , run : select
rds_manage_extension (' create ', ' oss_fdw ');
# Create a server instance
CREATE SERVER ossserver FOREIGN DATA WRAPPER oss_fdw
OPTIONS
( host ' oss - cn - hangzhou . aliyuncs . com ', id ' xxx
', key ' xxx ', bucket ' mybucket ');
# Create an OSS external table
CREATE FOREIGN TABLE ossexample
( date text, time text, open float,
high float, low float, volume int )
SERVER ossserver
OPTIONS ( filepath ' osstest / example . csv ', delimiter
', ' ,
format ' csv ', encoding ' utf8 ', PARSE_ERROR ' 100
');
# Create a table , to which data is loaded
create table example
( date text, time text, open float,
high float, low float, volume int );
# Load data from ossexample to example .
insert into example select * from ossexample ;
# As you can see
# oss_fdw estimates the file size on OSS and
formulates a query plan correctly .
explain insert into example select * from ossexample ;
QUERY PLAN

Insert on example ( cost = 0 . 00 .. 1 . 60 rows = 6 width
= 92 )
-> Foreign Scan on ossexample ( cost = 0 . 00 .. 1 . 60
rows = 6 width = 92 )
Foreign OssFile : osstest / example . csv . 0
Foreign OssFile Size : 728
( 4 rows )
# Write the data in the example table to OSS .
insert into ossexample select * from example ;
explain insert into ossexample select * from example ;
```

| QUERY | PLAN |
|---|------|
| Insert on ossexample (cost = 0 . 00 .. 16 . 60 rows = 660 width = 92) | |
| -> Seq Scan on example (cost = 0 . 00 .. 16 . 60 rows = 660 width = 92) | |
| (2 rows) | |

oss_fdw usage tips

- oss_fdw is an external table plugin developed based on the PostgreSQL FOREIGN TABLE framework.
- The data import performance is related to the PostgreSQL cluster resources (CPU I/O MEM MET) and OSS.
- For expected data import performance, ossendpoint in ossprotocol must match the region where PostgreSQL is located in Alibaba Cloud. For more information, see the reference links at the end of this document.
- If the error "oss endpoint userendpoint not in aliyun white list" is triggered during reading of SQL statements for external tables, use these [endpoints](#). If the problem persists, submit a trouble ticket.

Error handling

When an import or export error occurs, the error log contains the following information:

- code: HTTP status code of the erroneous request.
- error_code: Error code returned by OSS.
- error_msg: Error message provided by OSS.
- req_id: UUID that identifies the request. If you cannot solve the problem, you can seek help from OSS development engineers by providing the req_id.

For more information about error types, see the reference links at the end of this document. Timeout errors can be handled using oss_ext parameters.

- [OSS help](#)
- [PostgreSQL CREATE FOREIGN TABLE](#)
- [Exception handling](#)
- [OSS error response](#)

Hide ID and key

If ID and key parameters for CREATE SERVER are not encrypted, plaintext information is displayed using `select * from pg_foreign_server`, making

the ID and key exposed. The symmetric encryption can be performed to hide the ID and key (use different keys of different instances for further protection of your information). However, to avoid incompatibility with old instances, you cannot use methods similar to GP to add a data type.

Encrypted information:

```
postgres=# select * from pg_foreign_server ;
 srvname | srvowner | srvfdw | srvtype | srvversion |
-----+-----+-----+-----+-----+-----
 ossserver |      10 |  16390 |          |             |
 | { host = oss - cn - hangzhou - zmf . aliyuncs . com , id =
 MD5xxxxxxx x , key = MD5xxxxxxx x , bucket = 067862 }
```

The encrypted information is preceded by MD5 (total length: len%8==3). Therefore , encryption is not performed again when the exported data is imported. But you cannot create the key and ID preceded by MD5.