Alibaba Cloud ApsaraDB for MySQL

Quick Start for PPAS

Issue: 20190317

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.
`Courier font`	It is used for commands.	Run the `cd / d C :/ windows` command to enter the Windows system folder.
Italics	It is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	It indicates that it is a optional value, and only one item can be selected.	`ipconfig` *[-all|-t]*

Style	Description	Example
{} or {a|b}	It indicates that it is a required value, and only one item can be selected.	`swich {stand | slave}`

Contents

1 Limits

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

Operation	Description
Modify database parameter settings	Currently not supported
Database root permission	RDS does not offer the superuser permission to users.
Database backup	You can back up data only through pg_dump.
Data migration to the cloud	You can only use psql to restore data backed up by pg_dump.
Set up database replication	· You do not need to set up data replication because the system has automatically set up PPAS stream replication based the HA mode. · The PPAS slave node is invisible to users, and cannot be used directly for access.
Restart an RDS instance	You must restart an instance through the *RDS console* or APIs.
Network settings	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 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.

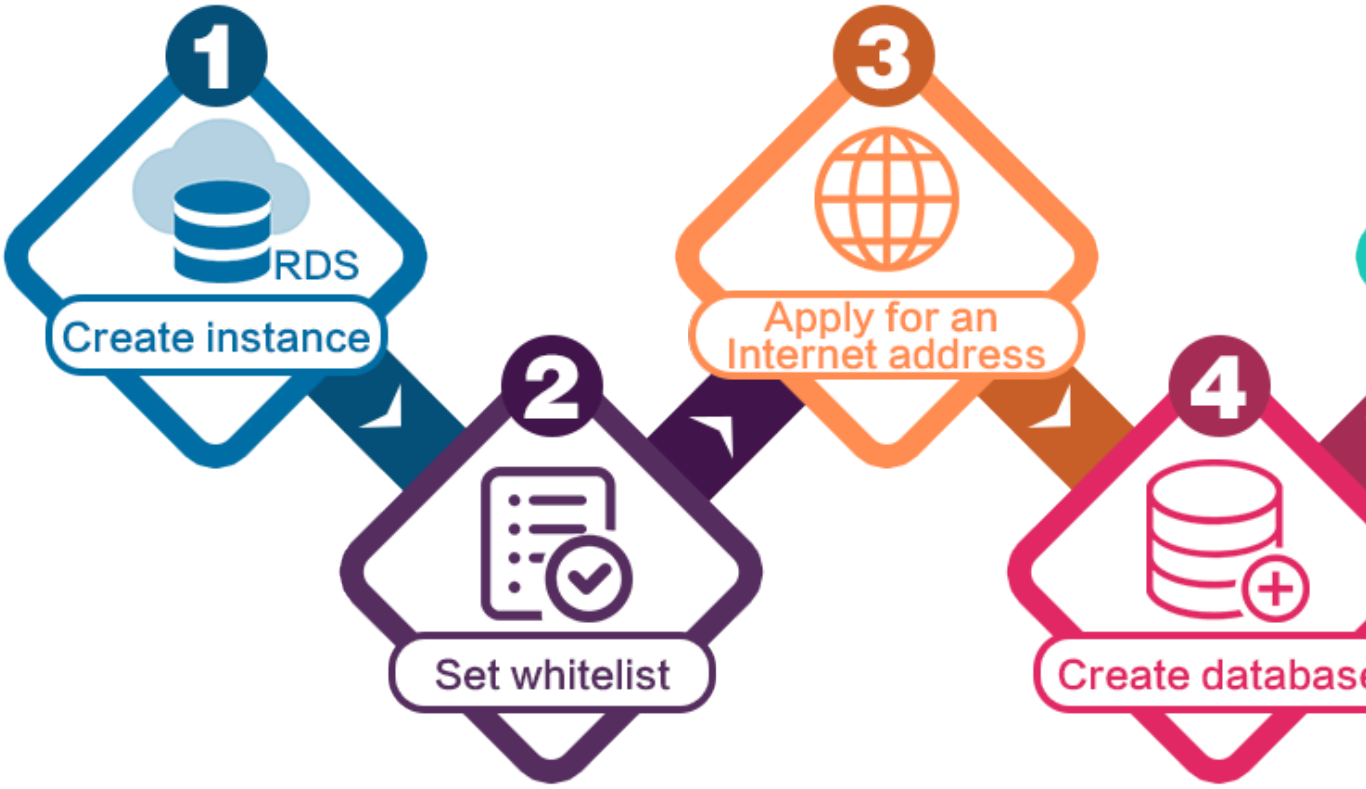

3 Create an instance

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 the APIs to create an instance, see *CreateDBInstance*.

Prerequisites

- You must have registered an Alibaba Cloud account.

Procedure

1. Log on to the *RDS console*.
2. On the Instances page, click Create Instance to open the Create page.
3. Select Subscription or Pay-As-You-Go. For more information about billing methods, 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 both 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). VPC needs to be created beforehand. Alternatively, you can change the network type after creating an instance. For more information, see *Set network type*.
- 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 type of instance and click Batch purchase in Purchase List.

6. Select Product Terms of Service and Service Level Notice and Terms of Use, and then:

   - Click Pay Now if the billing method of the instance is Subscription.
   - Click Activate if the billing method of the instance is Pay-As-You-Go.

4 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

You can adopt this method when using other clients. When you connect to an RDS instance through clients, choose to use the *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.

Procedure

1. Add the IP address that is allowed to access the RDS instance to the RDS whitelist. For more information, see *Set a 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 the Create - Server window, enter the server name, as shown in the following figure.

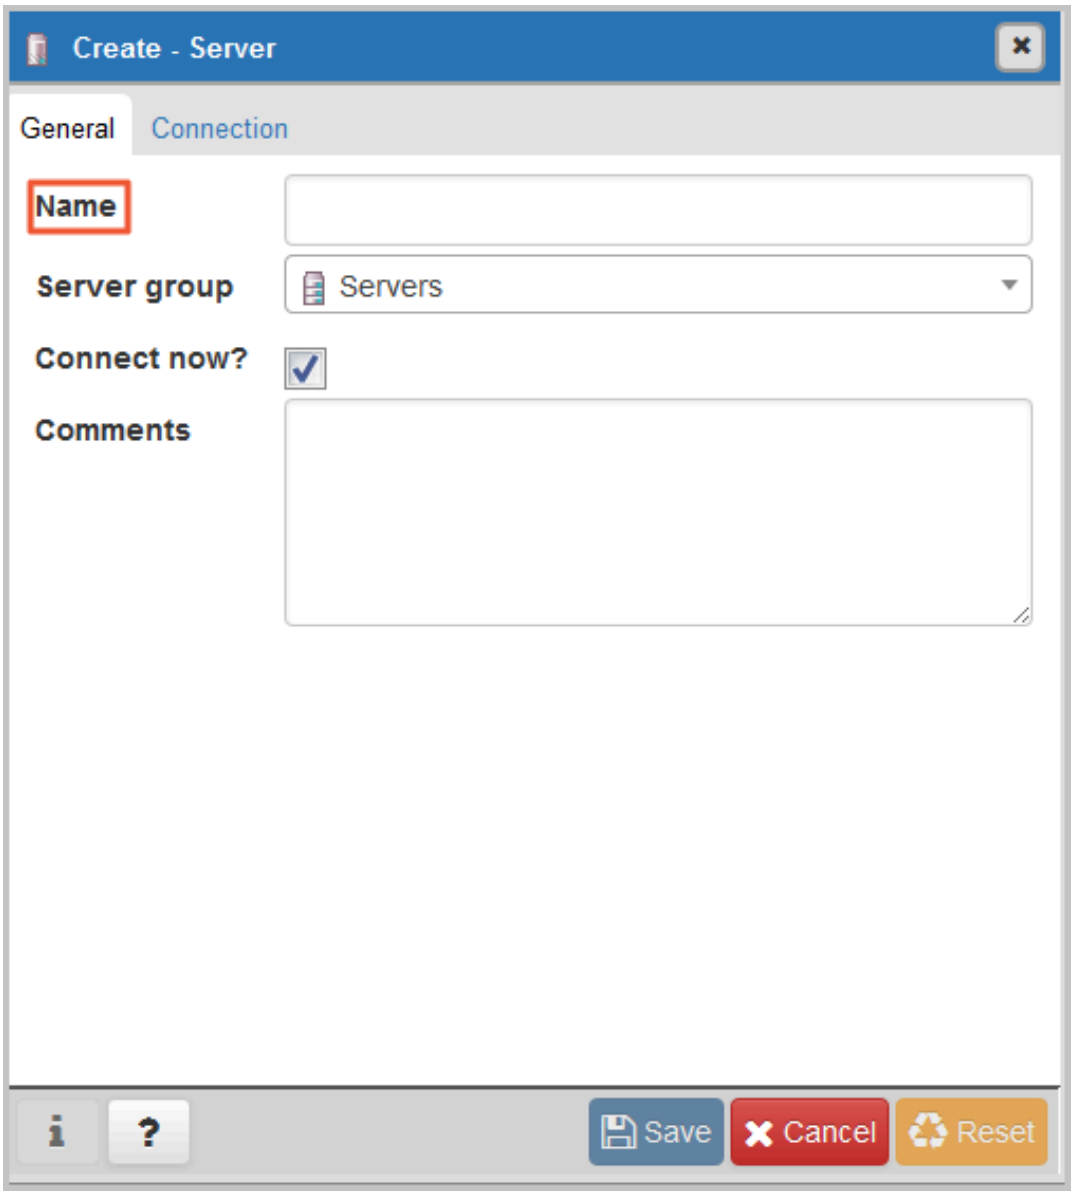

5. Click the Connection tab and enter information about the instance to be connected, as shown in the following figure.

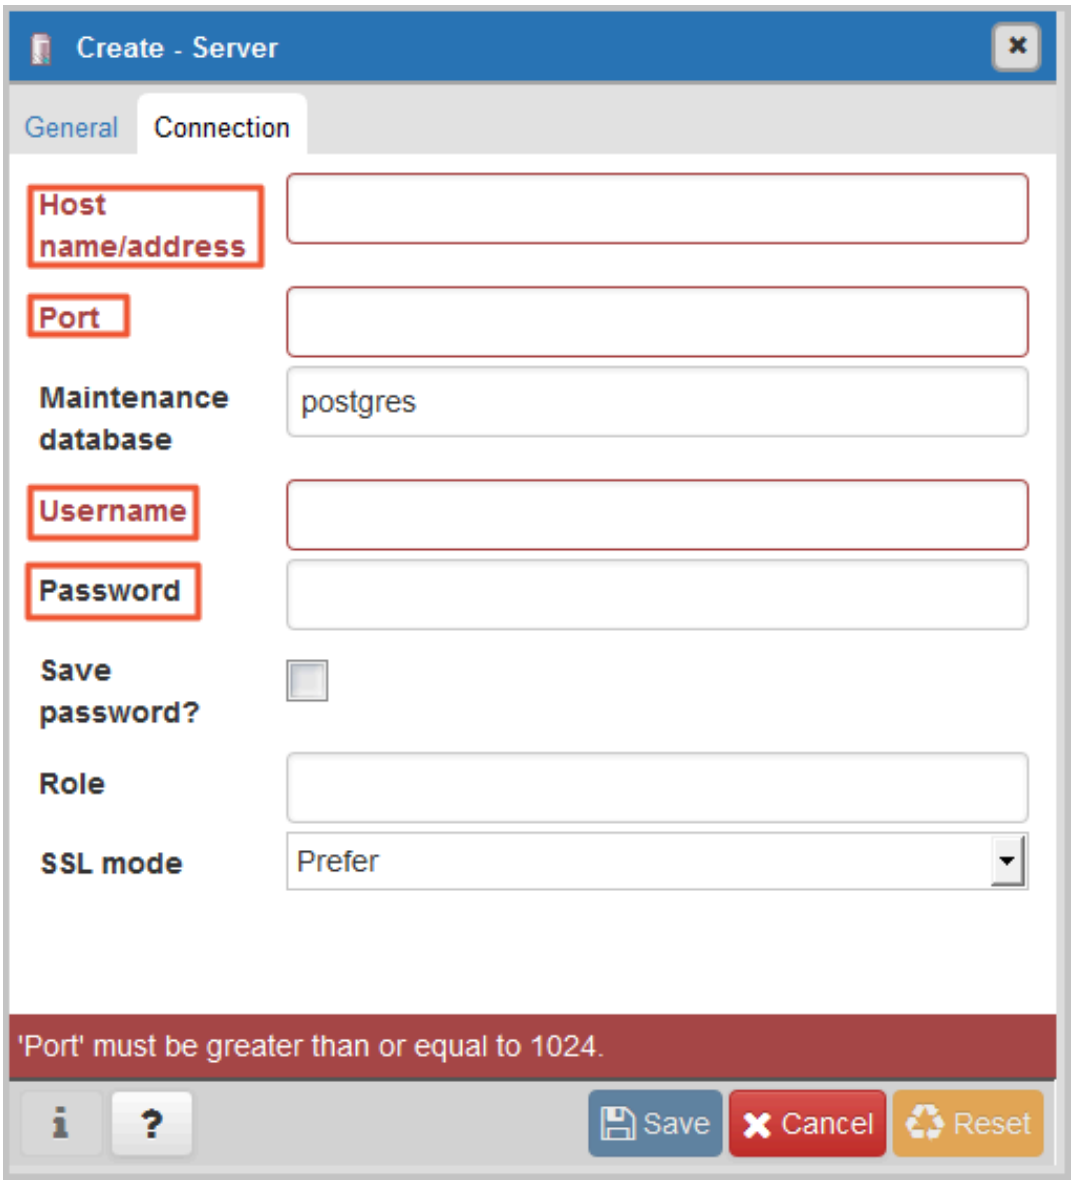

Parameter description:

- `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. Do 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. 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 > edb or postgres. The following interface is displayed, which indicates that the connection to RDS instance is successful.

Note:

edb and postgres are default system databases of the RDS instance. Do not perform any operation in these databases.

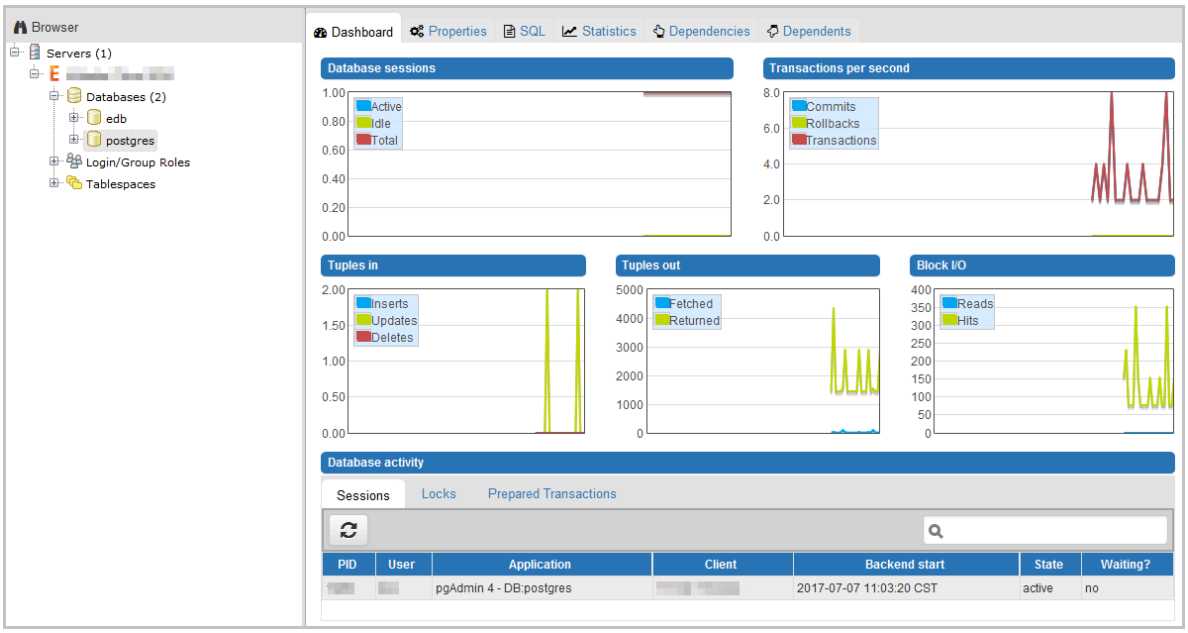

5 Read/write external data files using oss_fdw

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 transmissi on 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_flush_block_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_li  st_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_o  ne_file  = ' oss_test / test . gz .  2 ';
 select  *  from   oss_fdw_li  st_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_ERRO  RS  ' 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 |
srvacl |
                  srvoptions

-----------+----------+--------+---------+------------+--------
+-----------------------------------------------------------------------------------
-----------------------------------
  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.