

Alibaba Cloud

AnalyticDB for PostgreSQL FAQ

Document Version: 20220701

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 because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should 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 this document based on the "status quo", "being defective", and "existing functions" of its products and services. 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 take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, 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 contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, 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 this document 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 directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
 Danger	A danger notice 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.
 Warning	A warning notice 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 restart an instance.
 Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: If the weight is set to 0, the server no longer receives new requests.
 Note	A note indicates supplemental instructions, best practices, tips, and other content.	 Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings > Network > Set network type .
Bold	Bold formatting is used for buttons, menus, page names, and other UI elements.	Click OK .
<code>Courier font</code>	Courier font is used for commands	Run the <code>cd /d C:/window</code> command to enter the Windows system folder.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <i>Instance_ID</i>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

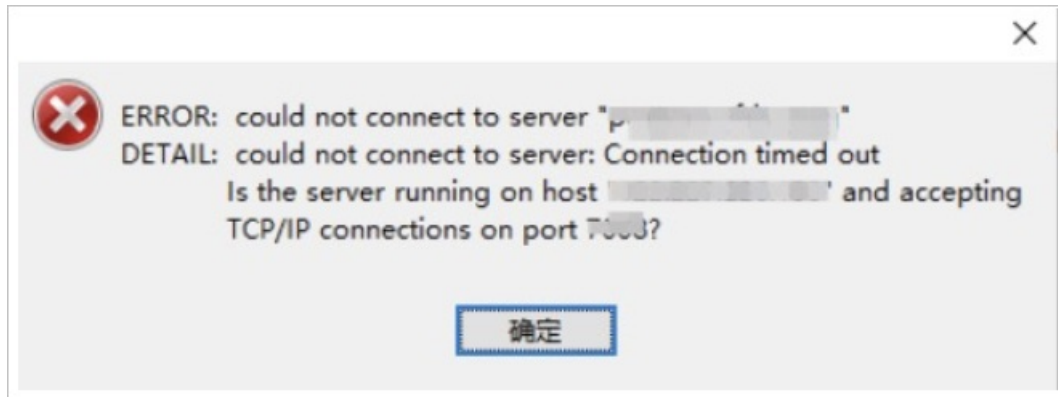
Table of Contents

1.What do I do if I cannot connect to an instance?	05
2.How to view the data size of a table or a database?	07
3.How do I obtain technical details about AnalyticDB for Postgr... ..	08
4.How do I choose an AnalyticDB for PostgreSQL instance type?	09
5.How do I switch the network type?	10
6.How do I clear data on a locked disk?	11
7.How do I reconfigure parameters?	12
8.How to enable column-store and compression in AnalyticDB fo... ..	13
9.What do I do if an AnalyticDB for PostgreSQL instance fails to... ..	14
10.Do I need to execute COPY statements as the superuser user?	15
11.Why am I unable to view the bills for my pay-as-you-go inst... ..	16
12.How do I terminate the abnormal execution of an SQL state... ..	17
13.How do I view the distribution key and partitions of a table?	18
14.How do I join the AnalyticDB for PostgreSQL technical suppo... ..	19
15.Mappings between compute node types and compute group	20
16.Obtain an AccessKey pair	21
17.How do I set the timeout period of SQL statements?	24

1. What do I do if I cannot connect to an instance?

Problem description

An error occurs when you connect to an AnalyticDB for PostgreSQL instance from a client, as shown in the following figure.



Causes

The AnalyticDB for PostgreSQL instance cannot communicate with the IP address in the error message. The following lists describe the possible causes:

- The IP address in the error message is a local area network (LAN) IP address.
- The IP address in the error message is not added to an IP address whitelist of the AnalyticDB for PostgreSQL instance.

Solution

Obtain the IP address of your on-premises client and add the IP address to an IP address whitelist of the AnalyticDB for PostgreSQL instance.

In complex network environments, you may be unable to find the IP address of your on-premises client or add the IP address to the an IP address whitelist. This section describes how to obtain the IP address of your on-premises client.

Procedure

1. Perform the following operations to add `0.0.0.0/0` to an IP address whitelist of the AnalyticDB for PostgreSQL instance:
 - i. Log on to the [AnalyticDB for PostgreSQL console](#).
 - ii. In the upper-left corner, select the region where the instance resides.
 - iii. Find the instance and click its ID. The **Basic Information** page appears.
 - iv. In the left-side navigation pane, click **Security Controls**. The **Security Controls** page appears.
 - v. On the **Whitelist Settings** tab, click **Modify** to the right of the default IP address whitelist. The **Modify Whitelist** panel appears.

- vi. In the **IP Addresses** field, delete the IP address 127.0.0.1 and enter `0.0.0.0/0` .

Note `0.0.0.0/0` indicates that all IP addresses are allowed to access the instance. This may raise security risks. We recommend that you delete 0.0.0.0/0 after you no longer need it.

- vii. Click **OK**.

2. Run the following command by using the psql tool to connect to the AnalyticDB for PostgreSQL instance. For more information about how to download the psql tool, see [Use client tools to connect to an instance](#).

```
psql -h yourgpdbaddress.gpdb.rds.aliyuncs.com -p 3432 -d postgres -U gpdbaccount
```

Parameter description:

- o -h: the host address.
 - o -p: the port used to connect to the database.
 - o -d: the name of the database. The default value is postgres.
 - o -U: the account used to connect to the database.
 - o You can run the `psql --help` command to view more options. You can also run the `\?` command to view the commands supported in psql.
3. After you connect to the database, run the following command in the SQL command line window to query the IP address of your on-premises client:

```
select * from pg_stat_activity;
```

The value of the CLIENT_ADDR field in the query results is the IP address of your on-premises client.

```
postgres=> select * from pg_stat_activity;
 datid | datname | procpid | sess_id | usesysid | username |          current_query          | waiting | client_addr | client_port | application_name | backend_start | xact_start |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
 10902 | postgres | 45909 | 254025 | 17239 | test123 | select * from pg_stat_activity; | f       | 100.11.24.5 | 31964 | psql              | 2019-02-19 14:07:00.37242+08 | 2019-02-19 14:06:26.170467+08 |
(1 row)
```

4. In the AnalyticDB for PostgreSQL console, delete `0.0.0.0/0` from the default IP address whitelist, and enter the IP address of your on-premises client to access the AnalyticDB for PostgreSQL instance.

2. How to view the data size of a table or a database?

Suppose that the table schema is `<schemaname>` and the table name is `<tablename>` .

- Run the following command to query the **total size of a table** (unit: MB, including the table index and data):

```
select pg_size_pretty(pg_total_relation_size('<schemaname>.<tablename>'));
```

- Run the following command to query the **data size of a table** (unit: MB, excluding the index):

```
select pg_size_pretty(pg_relation_size('<schemaname>.<tablename>'));
```

- Run the following command to query the **total size of all the partitions** in a partition table (unit: MB, including the table index and data):

```
select schemaname,tablename,round(sum(pg_total_relation_size(schemaname || '.' || partitiontablename))/1024/1024) "MB" from pg_partitions where schemaname='<schemaname>' and tablename='<tablename>' group by 1,2;
```

- Run the following command to query the **total size of all the tables under a schema** (unit: MB, including the index and data):

```
select schemaname ,round(sum(pg_total_relation_size(schemaname||'.'||tablename))/1024/1024) "Size_MB" from pg_tables where schemaname='<schemaname>' group by 1;
```

- Run the following command to query **the size of each database** (unit: MB):

```
select datname,pg_size_pretty(pg_database_size(datname)) from pg_database;
```

3. How do I obtain technical details about AnalyticDB for PostgreSQL?

This topic describes how to obtain technical details about AnalyticDB for PostgreSQL.

1. .
2. Visit the [Greenplum Database website](#) to read the Greenplum Database documentation.
3. Join the AnalyticDB for PostgreSQL technical support group on DingTalk (group No.: 11700737). In this group, you can receive online professional assistance and notifications on new features and promotions.

4. How do I choose an AnalyticDB for PostgreSQL instance type?

For more information, see [Instance types](#).

5. How do I switch the network type?


By default, the network type of is Virtual Private Cloud (VPC). If you want to switch the network type of an instance from the classic network to VPC, the instance and the VPC must be in the same region.

Background information

VPC: You can use a VPC to build an isolated network environment in Alibaba Cloud. You can customize route tables, CIDR blocks, and gateways in a VPC. If you want to migrate your applications to the cloud, you can build a virtual data center by connecting your self-managed data center to the resources in a VPC over a leased line or a virtual private network (VPN).

Procedure

1. Create a VPC in the same region with the AnalyticDB for PostgreSQL instance.
2. Log on to the [AnalyticDB for PostgreSQL console](#).
3. In the upper-left corner of the console, select the region where the instance resides.
4. Find the instance that you want to manage and click the instance ID.
5. In the left-side navigation pane, click **Database Connection**. The **Connection Information** page appears.
6. Click **Switch to VPC**. The **Switch to VPC** panel appears.
7. Select a VPC and a vSwitch. Click **OK**.

 **Notice** After you switch the network type of the instance from classic network to VPC, the network type of the internal endpoint is switched from classic network to VPC. ECS instances in the classic network cannot access the VPC-type AnalyticDB for PostgreSQL instances. The public endpoint remains unchanged.

API references

API	Description
ModifyDBInstanceNetworkType	Switches the network type of an instance between VPC and classic network.

6. How do I clear data on a locked disk?

Determine whether a disk is full

When the disk space of a compute node or coordinator node reaches its limit, the entire instance is locked.

Check whether an instance is locked

Connect to the instance and run the following command to check whether the instance is locked:

```
show rds_force_trans_ro_non_sup;
```

If the value of `rds_force_trans_ro_non_sup` is `on` in the response, the instance is locked and in the read-only state.

Clear data

When an instance is locked because the disk is full, TRUNCATE, DROP and GRANT operations are still supported on data tables. You can use these operations to clear data until the disk space falls within the threshold. Then, the instance is automatically unlocked within about five minutes.



Notice When the instance is locked, DELETE operations are not supported. When you execute DELETE statements, data is written to pg_xlog, which consumes disk space.

You can also execute the following statement to query the table size:

```
select pg_size_pretty(pg_total_relation_size('test'));
```

7. How do I reconfigure parameters?

You can reconfigure the parameters of a session used to establish a connection with your AnalyticDB for PostgreSQL instance. For more information, see the [Greenplum documentation](#).

The statement used to reconfigure a parameter is as follows: `ALTER ROLE <The username of your account> SET <The name of the parameter> = <The new value of the parameter>`. This statement only takes effect on the specified user.

If you want to reconfigure instance-level parameters, submit a ticket or join the AnalyticDB for PostgreSQL technical support group on DingTalk. For more information, see [How do I obtain technical details about AnalyticDB for PostgreSQL?](#)

8.How to enable column-store and compression in AnalyticDB for PostgreSQL?

Row-store with no compression is used by default when you create a table in AnalyticDB for PostgreSQL. To enable column-store and compression, you must specify the column-store and compression options when creating the table. For example, you can add the following statement to the tabulation statements to enable the two features.

```
with (APPENDONLY=true, ORIENTATION=column, COMPRESSTYPE=zlib, COMPRESSLEVEL=5, BLOCKSIZE=1048576, OIDS=false)
```

In general, we recommend that you enable column-store and compression, especially when you have a lot of complicated queries to the table, or when you want to reduce the storage cost. See [Define storage models for tables](#).

9. What do I do if an AnalyticDB for PostgreSQL instance fails to be created or remains in the creating state for a long time?

The creation of an AnalyticDB for PostgreSQL instance generally takes 20 minutes or less. If the AnalyticDB for PostgreSQL instance you purchased remains in the creating state for longer than 20 minutes, the zone you selected may not be able to provide sufficient resources. In this situation, you can [submit a ticket](#) to apply for more resources from the backend. Alternatively, you can select another zone and try again.

10. Do I need to execute COPY statements as the superuser user?

Symptom

I executed a `CREATE ROLE` statement to create a standard account named gpuser01 in AnalyticDB for PostgreSQL. However, when I executed a COPY statement to copy a script as the gpuser01 user, the system displayed " `ERROR: must be superuser to COPY to or from a file` ."

How do I promote a standard account to the superuser user?

Solution

- AnalyticDB for PostgreSQL does not support the superuser user. For more information, see [Features and limits](#).
- You can copy data to a file by executing:

```
psql -c 'copy xx to stdout' > file
```

Or

```
cat file | psql -c 'copy xx from stdin'
```

11. Why am I unable to view the bills for my pay-as-you-go instance? And when can I receive my invoice?

The bill for a pay-as-you-go instance within a month is generated on the second day of the next month. For example, the bill of February is only available since March 2. You can only receive the invoice for a month on the second day of the next month.

12. How do I terminate the abnormal execution of an SQL statement?

To terminate an SQL statement or session to restore the system status, use the `pg_stat_activity` system view to obtain the status of the current query:

```
select current_query,procpid from pg_stat_activity ;
           current_query                | procpid
-----+-----
select current_query,procpid from pg_stat_activity; | 32584
SELECT xxx                                         | 32238
```

Note `current_query` indicates the ongoing query and `procpid` indicates the ID of the background process.

Terminate the execution of a `SELECT` statement from a query other than the current query by executing:

```
SELECT pg_cancel_backend(pg_stat_activity.procpid)
FROM pg_stat_activity
WHERE procpid <> pg_backend_pid() and current_query like 'SELECT%'; -- Note that SELECT is entered in all uppercase.
pg_cancel_backend
-----
t
(1 rows)
```

Or

```
SELECT pg_terminate_backend(pg_stat_activity.procpid)
FROM pg_stat_activity
WHERE procpid <> pg_backend_pid() and current_query like 'SELECT%'; -- Note that SELECT is entered in all uppercase.
pg_terminate_backend
-----
t
(1 rows)
```

Note The `cancel` and `terminate` operations are only supported for queries initiated by the current user or by a user who has equal or fewer permissions than the current user. If the system displays "ERROR: must be superuser or rds_superuser to signal other server processes", queries initiated by other users are running in the background. This does not interrupt the execution of the `cancel` or `terminate` operations.

13.How do I view the distribution key and partitions of a table?

View the distribution key of a table

- Run the following command in psql to view the distribution key of a table:

```
\d <table_name>
```

- Execute the following SQL statement to view the distribution key of a table:

```
# AnalyticDB for PostgreSQL V4.3
SELECT attname FROM pg_attribute WHERE attrelid='<schema_name>.<table_name> '::regclass and
d attnum in (SELECT unnest(attrnums) FROM pg_catalog.gp_distribution_policy t WHERE local
oid='<schema_name>.<table_name> '::regclass);
# AnalyticDB for PostgreSQL V6.0
SELECT attname FROM pg_attribute WHERE attrelid='<schema_name>.<table_name> '::regclass and
d attnum in (SELECT unnest(distkey) FROM pg_catalog.gp_distribution_policy t WHERE localo
id='<schema_name>.<table_name> '::regclass);
```

Note

- `<schema_name>` : schema name of the table.
- `<table_name>` : The name of the table.

View the partitions of a table

- Run the following command in psql to view the partitions of a table:

```
\d+ <table_name>
```

- Execute the following SQL statement to view the partitions of a table:

```
SELECT pg_get_partition_def('<schema_name>.<table_name> '::regclass,true);
```

Note

- `<schema_name>` : schema name of the table.
- `<table_name>` : The name of the table.

14.How do I join the AnalyticDB for PostgreSQL technical support group on DingTalk?

You can join our DingTalk group by searching for 11700737. This DingTalk group presents online technical support from AnalyticDB for PostgreSQL experts, regularly published promotional offers, and new features to come.

15. Mappings between compute node types and compute group types

On August 23, 2019, the AnalyticDB for PostgreSQL instance type definition was changed. Instances have been changed from being composed of **compute groups** to being composed of **compute nodes**. This topic describes the mappings between **compute nodes** and **compute groups**.

Compute group is replaced with **compute node** in the definition to simplify the selection of MPP cluster types and comply with cluster database naming conventions. Each **compute node** only contains one MPP data partition, which simplifies the selection and use of MPP cluster types. A **compute group** may contain multiple MPP data partitions based on its type.

Compute group type (no longer used)	Compute node type	Number of MPP data partitions	Total amount of resources
1x2C SSD	Two 1-core SSD nodes	2	2-core/16 GB memory/160 GB user storage/320 GB total storage
1x16C SSD	Sixteen 1-core SSD nodes	16	16-core/128 GB memory/1,280 GB user storage/2,560 GB total storage
4x4C SSD	Four 4-core SSD nodes	4	16-core/128 GB memory/1,280 GB user storage/2,560 GB total storage
2x2C HDD	Two 2-core HDD nodes	2	4-core/32 GB memory/2 TB user storage/4 TB total storage
4x4C HDD	Four 4-core HDD nodes	4	16-core/96 GB memory/8 TB user storage/16 TB total storage

According to the preceding mappings, **compute nodes** and **compute groups** have the same performance, capacity, and price when they share the same amount of resources and same number of partitions.


16. Obtain an AccessKey pair

You can create an AccessKey pair for an Alibaba Cloud account or a RAM user. When you call API operations, you must use the AccessKey pair to complete identity verification.

Context

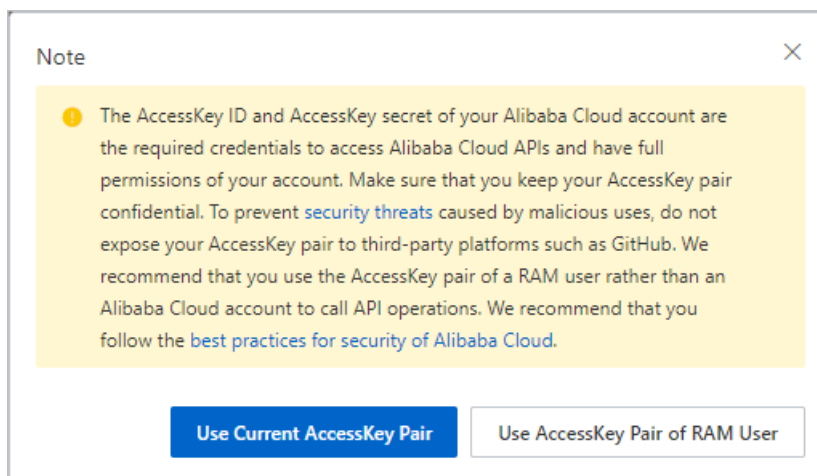
An AccessKey pair consists of an AccessKey ID and an AccessKey secret.

- The AccessKey ID is used to identify a user.
- The AccessKey secret is used to verify the identity of the user. You must keep your AccessKey secret strictly confidential.

 **Warning** If the AccessKey pair of your Alibaba Cloud account is leaked, your resources are exposed to potential risks. We recommend that you use the AccessKey pair of a RAM user to call API operations. This prevents the AccessKey pair of your Alibaba Cloud account from being leaked.

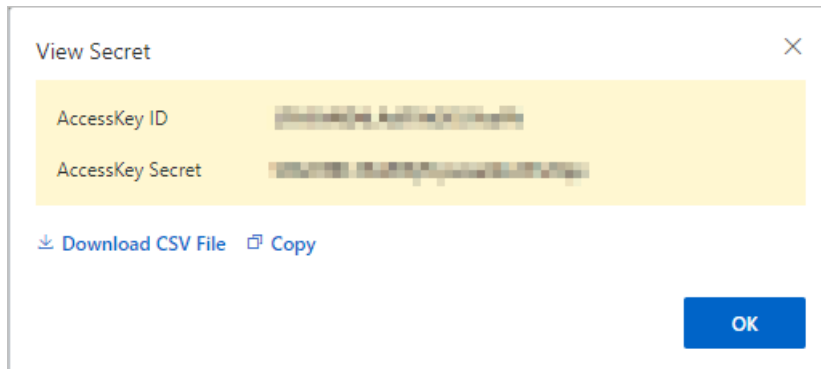
Procedure

1. Log on to the [Alibaba Cloud Management Console](#) with your Alibaba Cloud account.
2. Move the pointer over the profile picture in the upper-right corner of the page and click **AccessKey Management**.
3. In the **Note** dialog box, click **Use Current AccessKey Pair** or **Use AccessKey Pair of RAM User**.




- o Obtain the AccessKey pair of the Alibaba Cloud account
 - a. Click **Use Current AccessKey Pair**.
 - b. On the **AccessKey Pair** page, click **Create AccessKey**.

- c. In the **View Secret** dialog box, view the AccessKey ID and AccessKey secret. Click **Download CSV File** to download the AccessKey pair or click **Copy** to copy the AccessKey pair.

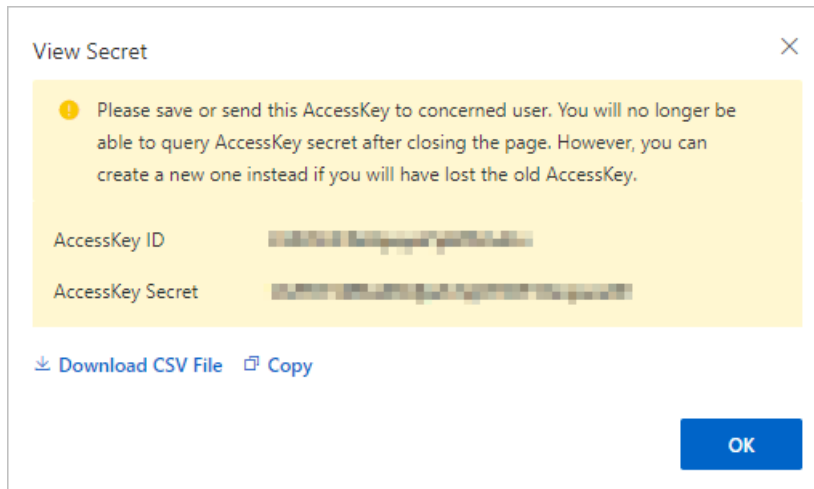


- o Obtain the AccessKey pair of a RAM user
 - a. Click **Use AccessKey Pair of RAM User**.
 - b. On the **Users** page of the **RAM console**, find the RAM user whose AccessKey pair you want to obtain.

 **Note** If you do not have a RAM user, create one first. For more information, see [Create a RAM user](#).

- c. Click the name of the RAM user in the User Logon Name/Display Name column.
- d. In the **User AccessKeys** section of the **Authentication** tab, click **Create AccessKey**.

- e. In the **View Secret** dialog box, view the AccessKey ID and AccessKey secret. Click **Download CSV File** to download the AccessKey pair or click **Copy** to copy the AccessKey pair.



Note

- An AccessKey secret is displayed only after you click Create AccessKey. You cannot query the AccessKey secret in subsequent operations. Therefore, you must back up your AccessKey secret.
- If your AccessKey pair is leaked or lost, you must create another AccessKey pair. You can create a maximum of two AccessKey pairs for each RAM user.

17. How do I set the timeout period of SQL statements?

This topic describes how to set the timeout period of SQL statements in .

Problem description

The following error message is returned after an SQL statement is executed for more than 3 hours:

```
ERROR: canceling statement due to statement timeoutTime: 10801445.540 ms (03:00:01.446)
```

Causes

By default, the timeout period of SQL statements in is set to 3 hours. If an SQL statement is executed for more than 3 hours, the execution of the statement is terminated.

Solutions

provides the `statement_timeout` parameter to modify the timeout period of SQL statements. You can use one of the following methods to modify the `statement_timeout` value:

- Perform the following steps in the AnalyticDB for PostgreSQL console (instance-level modification):

i.

ii.

iii.


iv. In the left-side navigation pane, click **Parameter Configuration**.

v. On the **Parameter Configuration** page, find the `statement_timeout` parameter and click the

 icon in the **Running Value** column.

vi. Enter a new parameter value.

- Valid values: 0 to 2147483647.
- If you set the value to 0, the timeout feature of SQL statements is disabled.
- Unit: milliseconds.

 **Note** We recommend that you set the timeout period of SQL statements to a value greater than 60000 milliseconds to prevent impacts on other tasks.

vii. Click **OK**.

viii. Click **Submit** in the upper-right corner.


ix. In the **Parameter Modification** message, click **OK**.

- Execute the following SQL statement (session-level modification):

```
SET statement_timeout = 0;
```

- Execute the following SQL statement (user-level modification):

```
ALTER ROLE <username> SET statement_timeout TO 600000;
```


 **Note** Replace `username` with the database account that is used to set the timeout period of SQL statements.