

Alibaba Cloud

AnalyticDB for PostgreSQL Quick Start

Document Version: 20220120

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. Get started with AnalyticDB for PostgreSQL	05
2. Create an AnalyticDB for PostgreSQL instance	06
3. Configure an IP address whitelist	11
4. Create a database account	14
5. Use client tools to connect to an instance	16

1. Get started with AnalyticDB for PostgreSQL

This is the AnalyticDB for PostgreSQL Quick Start. AnalyticDB for PostgreSQL is a distributed database service that provides massively parallel processing (MPP) capabilities. By completing the operations in this guide, you will learn how to use the console to create and manage AnalyticDB for PostgreSQL instances.

If you use AnalyticDB for PostgreSQL for the first time, we recommend that you read the following topics first:

- [Overview](#): provides an overview of AnalyticDB for PostgreSQL and describes its features and scenarios.
- [Pricing](#): describes the pricing and billing methods of AnalyticDB for PostgreSQL.
- [Instance specifications](#): describes how to select appropriate specifications for your AnalyticDB for PostgreSQL instances.

To use AnalyticDB for PostgreSQL, you must perform the following operations:

1. [Create an AnalyticDB for PostgreSQL instance.](#)
2. [Configure an IP address whitelist.](#)
3. [Create a database account.](#)
4. [Use client tools to connect to an instance.](#)

2. Create an AnalyticDB for PostgreSQL instance


This topic describes how to create an instance in the AnalyticDB for PostgreSQL console.

Prerequisites

- An Alibaba Cloud account is created. If you do not have an Alibaba Cloud account, visit the [Alibaba Cloud official website](#) to create an account.
- If this is the first time that you create an AnalyticDB for PostgreSQL instance, make sure that a service-linked role is created. For more information, see [Service linked role](#).
 - i. Log on to the [AnalyticDB for PostgreSQL console](#).
 - ii. In the upper-right corner, click **Create Instance**.
 - iii. In the **Create Service Linked Role** message, click **OK**.



Procedure




1. Log on to the [AnalyticDB for PostgreSQL console](#).
2. In the upper-right corner, click **Create Instance** to go to the buy page.
3. Configure the **Product Type** parameter. Valid values:
 - **Subscription** : You must pay a subscription fee when you create an instance. This billing method is more cost-effective than the pay-as-you-go billing method and is suitable for long-term use. Larger discounts are provided for longer subscription periods.
 - **Pay-as-you-go**: You are charged on an hourly basis for the resources that you use. The pay-as-you-go billing method is suitable for short-term use. If you no longer need a pay-as-you-go AnalyticDB for PostgreSQL instance, you can release the instance to reduce costs.




 **Note** You can change the billing method of an instance from pay-as-you-go to subscription. However, you cannot change the billing method from subscription to pay-as-you-go. For more information, see [Change the billing method from pay-as-you-go to subscription](#).

4. Configure the parameters that are described in the following table.

Parameter	Description
-----------	-------------


Parameter	Description
Region	<p>Select the region where you want to create the instance.</p> <div> <p> Note</p> <ul style="list-style-type: none"> After you complete the purchase, you cannot change the region of the instance. To accelerate access from users to the instance, we recommend that you select a region that is close to the geographical location of the users. Make sure that the instance is deployed in the same region as the Elastic Compute Service (ECS) instance to which you want to connect. Otherwise, the AnalyticDB for PostgreSQL instance and the ECS instance can communicate only over the Internet. As a result, the performance of the instance may be compromised. </div>
Zone	<p>A zone is an independent geographical location in a region. All zones in a region provide the same level of service performance.</p> <p>You can create the instance in the same zone as the ECS instance to which you want to connect or in a different zone.</p>
Network Type	Select the network type of the instance. You can select only VPC .
VPC	<p>Select the ID of the virtual private cloud (VPC).</p> <p>A VPC is an isolated virtual network that provides higher security and higher performance than the classic network. Before you configure this parameter, make sure that a VPC and a vSwitch are created in the region where you want to create the instance. For more information, see Create and manage a VPC.</p>
vSwitch	Select a vSwitch in the specified VPC.
Instance Resource Type	<p>Select the resource type of the instance. Valid values:</p> <ul style="list-style-type: none"> Serverless: You can specify only the required computing resources without the need to reserve storage resources. Flexible Storage Mode: You can separately expand disks and smoothly scale out instances online. Reserved Storage Mode: You cannot separately expand disks or smoothly scale out instances. <div> <p> Note The resource types that are available vary based on the region and zone. If the current region and zone do not provide the required instance resource type, change the region and zone.</p> </div>

Parameter	Description
Edition	<p>Select an edition. Valid values: High-availability and High Performance (Basic Edition).</p> <p> Note You can select High Performance (Basic Edition) only if you select Reserved Storage Mode for the Instance Resource Type parameter. For more information about Basic Edition, see Basic Edition.</p>
Coordinator Nodes	<p>Select the number of coordinator nodes. Default value: 1.</p> <p>The maximum number of coordinator nodes that you can select in the console is 2. If you require more coordinator nodes, Submit a ticket to contact technical support.</p>
Compute Node Specifications Node Cores	<p>Select specifications for the compute nodes. The available specifications vary based on the resource type and edition. For more information, see Instance specifications.</p> <p> Note During the public preview from January 01, 2022 to January 31, 2022, you can select only 4C16G for instances in Serverless mode.</p>
Nodes	<p>Select the number of compute nodes. The instance performance increases with the number of compute nodes. The number of compute nodes that you can select varies based on the resource type and edition:</p> <ul style="list-style-type: none"> ◦ Serverless <p>High-availability: Valid values: 2 to 512. The number must be an integral multiple of two.</p> <p> Note During the public preview from January 01, 2022 to January 31, 2022, you can select up to 12 compute nodes for instances in Serverless mode.</p> ◦ Elastic Storage Mode <ul style="list-style-type: none"> ■ High-availability: Valid values: 4 to 512. The number must be an integral multiple of four. ■ High Performance (Basic Edition): Valid values: 2 to 512. The number must be an integral multiple of two. ◦ Reserved Storage Mode <p>High-availability: Valid values: 2, 4, 8, 12, 16, 24, 32, 64, 96, and 128.</p>

Parameter	Description
Storage Disk Type Storage Type	<p>Select a disk type. The available disk types vary based on the resource type and edition:</p> <ul style="list-style-type: none"> ◦ Serverless: You do not need to select a disk type. ◦ Elastic Storage Mode <ul style="list-style-type: none"> ▪ High-availability: You can select Enhanced SSD (ESSD) or Ultra Disk. ▪ High Performance (Basic Edition): You can select only Enhanced SSD (ESSD). ◦ Reserved Storage Mode: You can select only SSD.
Encryption Type	<p>Select Not Encrypted or Disk Encryption.</p> <p> Note You can configure this parameter only if you select Elastic Storage Mode. Disk encryption is not supported for instances in Serverless mode or reserved storage mode.</p>
CMK	<p>Select a disk encryption key.</p> <p>This parameter is required only if you select Elastic Storage Mode and Disk Encryption.</p>
Single Node Storage Capacity	<p>Select the exclusive storage capacity for each node in the instance. Valid values: 50 to 4,000. Unit: GB. The value must be an integral multiple of 50.</p> <p> Note You can configure this parameter only if you select Elastic Storage Mode.</p>
Purchase Plan	<p>Select a subscription duration for the instance.</p> <p> Note This parameter is required only if you select Subscription.</p>

- Click **Buy Now**.
- On the **Confirm Order** page, perform the following operations based on the value of the **Product Type** parameter:
 - **Subscription**
 - Confirm information in the **Parameters** section, select **Terms of Service**, and then click **Pay**.
 - On the **Purchase** page, click **Subscribe**.
 - **Pay-as-you-go**

On the **Confirm Order** page, confirm the information in the **Parameters** section, select **Terms of Service**, and then click **Activate Now**.
- After the instance is created, you can view the instance on the **Instances** page.

 **Note** It requires a specific amount of time to initialize an AnalyticDB for PostgreSQL instance. You can perform operations only when the instance is in the **Running** state.

Related API operations

API	Description
CreateDBInstance	Creates an AnalyticDB for PostgreSQL instance. You can call this operation to create an AnalyticDB for PostgreSQL instance in Serverless mode, elastic storage mode, or reserved storage mode.
CreateECSDBInstance	Creates an AnalyticDB for PostgreSQL instance in elastic storage mode.

3.Configure an IP address whitelist

To ensure the security and stability of AnalyticDB for PostgreSQL databases, AnalyticDB for PostgreSQL instances block access from all IP addresses by default. Before you use an AnalyticDB for PostgreSQL instance, you must add IP addresses or CIDR blocks that are used to access the AnalyticDB for PostgreSQL instance to the whitelists of the instance. A properly configured IP address whitelist can make your AnalyticDB for PostgreSQL instance more secure. We recommend that you maintain IP address whitelists on a regular basis.

Prerequisites

Before you configure a whitelist for an AnalyticDB for PostgreSQL instance, you must obtain the IP address of the client based on its installation location by using the following methods.

- ECS instances:
 - Linux instance
Run the `ifconfig` command to view NIC information. You can view the IP addresses, subnet masks, gateways, DNS servers, and MAC addresses in the command output.
 - Windows instance
In Command Prompt, run the `ipconfig /all` command to view NIC information. You can view the IP addresses, subnet masks, gateways, DNS servers, and MAC addresses in the command output.
- On-premises devices or third-party clouds:
 - Linux operating system: Run the `curl ipinfo.io |grep ip` command on the on-premises device to obtain its public IP address.
 - Windows operating system: Visit [ipinfo](#) on the on-premises device to obtain its public IP address.

Procedure

1. Log on to the [AnalyticDB for PostgreSQL console](#).
2. In the upper-left corner of the console, select the region where the instance resides.
3. Find the instance that you want to manage and click its ID.
4. In the left-side navigation pane, click **Security Controls**.
5. On the **Security Controls** page, perform the following operations:
 - Create a whitelist.
 - a. Click **Create Whitelist**.

b. In the **Create Whitelist** dialog box, specify the following parameters.


Parameter	Description
Whitelist Name	<p>The name of the new whitelist.</p> <ul style="list-style-type: none">■ The name can contain lowercase letters, digits, and underscores (_).■ The name must start with a lowercase letter and end with a lowercase letter or a digit.■ The name must be 2 to 32 characters in length.
IP Addresses	<p>The IP addresses or CIDR blocks that are allowed to access the instance.</p> <ul style="list-style-type: none">■ Separate multiple IP addresses with commas (.). A maximum of 999 unique IP addresses can be specified.■ Supported formats are specific IP addresses such as 10.23.12.24 and CIDR blocks such as 10.23.12.24/24. /24 indicates the length of the IP address prefix. An IP address prefix can be 1 to 32 bits in length.■ If you set the prefix length to 0, for example, 0.0.0.0/0 or 127.0.0.1/0, all IP addresses are allowed to access the instance. This poses a high security risk. Proceed with caution.■ The IP address 127.0.0.1 indicates that no external IP addresses are allowed to access the instance.

c. Click **OK**.

o Modify a whitelist.

a. Click **Modify** to the right of the whitelist name.

b. Add or remove IP addresses or CIDR blocks in the **IP Addresses** section.

 **Note** The **Whitelist Name** of the default whitelist cannot be modified.

c. Click **OK**.

o Delete a whitelist.

 **Note** The default whitelist cannot be deleted.

a. Click **Delete** to the right of the whitelist name.

b. In the **Delete Whitelist** message, click **OK**.

o Clear the default whitelist.

a. Click **Clear** to the right of the default whitelist.

b. In the **Clear Whitelist** message, click **OK**.

The default whitelist contains only 127.0.0.1 after it is cleared.

Related operations

Operation	Description
DescribeDBInstanceIPArrayList	Queries the IP addresses that are allowed to access an instance.
ModifySecurityIps	Modifies an IP address whitelist of an instance.


4. Create a database account

This topic describes how to create a database account for an AnalyticDB for PostgreSQL instance.

Context

AnalyticDB for PostgreSQL provides the following types of database accounts:

- **Privileged account:** an account that has all permissions on all databases. The first account created for an instance in the console is a privileged instance.
- **Standard accounts:** accounts that have all permissions only on their authorized databases.

 **Note** Permissions include SELECT, INSERT, UPDATE, DELETE, TRUNCATE, REFERENCES, and TRIGGER. For information about how to configure permissions, see [Manage users and permissions](#).

Precautions

- Before you use databases in an AnalyticDB for PostgreSQL instance, you must create a privileged account.
- Standard accounts cannot be created by using the console. For information about how to create standard accounts, see [Execute SQL statements to create accounts](#).
- After you create a privileged account for an instance, you cannot delete the privileged account.


Create a privileged account

1. Log on to the [AnalyticDB for PostgreSQL console](#).
2. In the upper-left corner of the console, select the region where the instance resides.
3. Find the instance that you want to manage and click its ID.
4. In the left-side navigation pane, click **Account Management**.
5. Click **Create Account**.
6. In the **Create Account** panel, specify the following parameters.

Parameter	Description
Account	<p>The name of the privileged account.</p> <ul style="list-style-type: none">◦ The name can contain lowercase letters, digits, and underscores (_).◦ The name must start with a lowercase letter and end with a lowercase letter or a digit.◦ The name cannot start with gp.◦ The name must be 2 to 16 characters in length.

Parameter	Description
New Password	<p>The password of the privileged account.</p> <ul style="list-style-type: none">◦ The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.◦ The following special characters are supported: ! @ # \$ % ^ & * () _ + - =◦ The password must be 8 to 32 characters in length.
Confirm Password	Enter the password of the account again.

7. Click **OK**.

 **Notice** After the account is created, you can click **Reset Password** in the **Actions** column to modify the account password. To ensure data security, we recommend that you change your password on a regular basis and do not use passwords that you have used before.

Execute SQL statements to create accounts

- Create a privileged account.

```
CREATE ROLE <Account name> WITH LOGIN ENCRYPTED PASSWORD <Password> RDS_SUPERUSER;
```

Example:

```
CREATE role admin0 WITH LOGIN ENCRYPTED PASSWORD '111111' rds_superuser;
```

- Create a standard account.

```
CREATE ROLE <Account name> WITH LOGIN ENCRYPTED PASSWORD <'Password'>;
```

Example:

```
CREATE role test1 WITH LOGIN ENCRYPTED PASSWORD '111111';
```

Related operations

Operation	Description
CreateAccount	Creates a privileged account.
DescribeAccounts	Queries the accounts created on an AnalyticDB for PostgreSQL instance.
ModifyAccountDescription	Modifies the name of an account for an instance.
ResetAccountPassword	Resets the password of an account for an instance.


5. Use client tools to connect to an instance

This topic describes how to connect to an AnalyticDB for PostgreSQL instance by using tools that support the message-based protocol of PostgreSQL.

Background information

AnalyticDB for PostgreSQL is fully compatible with the message-based protocol of PostgreSQL. You can connect to an AnalyticDB for PostgreSQL instance by using tools that support the protocol and GUI tools.

- Tools that support the message-based protocol of PostgreSQL include psql, libpq, Java Database Connectivity (JDBC), Open Database Connectivity (ODBC), and psycopg2.
- GUI tools include Alibaba Cloud Data Management (DMS) and DBeaver 6.0.0 or later.

 **Note** AnalyticDB for PostgreSQL V4.3 is built based on PostgreSQL 8.3. AnalyticDB for PostgreSQL V6.0 is built based on PostgreSQL 9.4.

DMS

DMS allows you to manage relational, online transaction processing (OLTP), online analytical processing (OLAP), and NoSQL database services. Relational database services include ApsaraDB RDS for MySQL, ApsaraDB RDS for SQL Server, ApsaraDB RDS for PostgreSQL, and ApsaraDB RDS for PPAS. OLTP database services include Distributed Relational Database Service (DRDS). OLAP database services include AnalyticDB for PostgreSQL and Data Lake Analytics (DLA). NoSQL database services include ApsaraDB for MongoDB and ApsaraDB for Redis. DMS offers an integrated solution to manage data, schemas, and servers. You can also use DMS to authorize users, audit security, view BI charts and data trends, track data, and optimize performance.

This section describes how to use DMS to log on to an AnalyticDB for PostgreSQL instance.

1. Log on to the [AnalyticDB for PostgreSQL console](#).
2. (Optional) Create an AnalyticDB for PostgreSQL instance. For more information, see [Create an AnalyticDB for PostgreSQL instance](#).

If an AnalyticDB for PostgreSQL instance has already been created, skip this step.

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. (Optional) Create a database account. For more information, see [Create a database account](#).

If an initial account has already been created, you can use this account to log on to the instance.

6. In the upper-right corner, click **Login Database**.
7. In the **Login instance** dialog box, specify the **Database Account** and **Database password** parameters, and click **Login**.

 **Note** AnalyticDB for PostgreSQL

For more information about DMS, see [Overview](#).

psql

psql is a common command-line tool used together with Greenplum and provides a variety of commands. Its binary files are located in the bin directory of Greenplum.

Download client tools:

For Red Hat Enterprise Linux (RHEL) 6 or 7 and CentOS 6 or 7, you can download client tools from the following links and decompress the tool package:

- Download links of client tools for AnalyticDB for PostgreSQL V4.3:
 - RHEL 6 or CentOS 6: [ADBPg_client_package_el6](#)
 - RHEL 7 or CentOS 7: [ADBPg_client_package_el7](#)
- Download links of client tools for AnalyticDB for PostgreSQL V6.0:
 - RHEL 6 or CentOS 6: [ADBPg_client_package_el6](#)
 - RHEL 7 or CentOS 7: [ADBPg_client_package_el7](#)

After you download the client tool package, you must decompress the package and install the tools. This section describes how to install the client tools for AnalyticDB for PostgreSQL V6.0 on CentOS 7.

1. Run the following command to decompress the client tool package in its directory:

```
tar -xzf adbpkg_client_package.el7.x86_64.tar.gz
```

2. Switch to the *bin* directory and run the following command:

```
cd adbpkg_client_package/bin
```

3. The *bin* directory includes client tools such as psql and pg_dump. Run the corresponding commands based on the reference documentation of each tool.
 - For more information about how to use psql to connect to an AnalyticDB for PostgreSQL instance, see [psql](#).
 - pg_dump is a logical backup tool for PostgreSQL. For more information about how to use pg_dump, see [pg_dump](#).

To use the psql tool to connect to an instance, perform the following steps:

1. Use one of the following methods to connect to an instance:
 - Connection string


```
psql "host=yourgpdbaddress.gpdb.rds.aliyuncs.com port=5432 dbname=postgres user=gpdbackcount password=gpdbackpassword"
```

- Specified parameters

```
psql -h yourgpdbaddress.gpdb.rds.aliyuncs.com -p 5432 -d postgres -U gpdbaccount
```

Parameter description:

- -h: the host address.
- -p: the port number used to connect to the instance.
- -d: the name of the database. The default value is postgres.
- -U: the account used to connect to the instance.

 **Note** You can run the `psql --help` command to view more options. You can also run the `\?` command to view the commands supported in psql.

2. Enter the password to go to the psql command interface.

References:

- For more information about how to use psql in Greenplum, see [Greenplum psql](#).
- AnalyticDB for PostgreSQL also supports psql for PostgreSQL. Take note of the differences of psql commands between Greenplum and PostgreSQL. For more information, see [PostgreSQL psql](#).

You can use the Docker toolbox for AnalyticDB for PostgreSQL, as shown in the following sample code. For more information about how to install Docker, visit the [Docker official website](#).

- Run the tool image corresponding to AnalyticDB for PostgreSQL V4.3.

```
docker run -idt --name=adbpgcli aliadbpg/adbpgcli:v4.3.0
docker exec -it adbpgcli /bin/bash -l
```


- Run the tool image corresponding to AnalyticDB for PostgreSQL V6.0.

```
docker run -idt --name=adbpgcli aliadbpg/adbpgcli:v6.3.0
docker exec -it adbpgcli /bin/bash -l
```

DBeaver

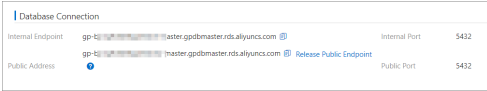
DBeaver can run on major operating systems such as Windows, macOS, and Linux. This section describes how to use DBeaver on Windows to connect to an AnalyticDB for PostgreSQL instance.

1. Download and run the installation package of DBeaver Community Edition. For more information about the download link, visit the [DBeaver official website](#).
2. Start DBeaver and choose **Database > New Database Connection**.
3. In the **Connect to a database** dialog box, select the **Greenplum** database type and click **Next**.

 **Note** The first time you connect to an instance, you must download the corresponding database driver file.

4. On the **Main** tab of the **Connection Settings** page, set the following parameters.

Parameter	Description	Example
-----------	-------------	---------

Parameter	Description	Example
Host	The endpoint and port number used to connect to the AnalyticDB for PostgreSQL instance. To obtain this information, perform the following steps:	gp-bp1g*****-master.gpdbmaster.rds.aliyuncs.com
Port	<ol style="list-style-type: none"> Log on to the AnalyticDB for PostgreSQL console. In the upper-left corner of the console, select the region where the instance resides. Find the instance that you want to manage and click the instance ID. On the Basic Information page, view Internal Endpoint and Internal Port or Public Address and Public Port in the Database Connection section. 	5432
Database	The name of the database.	postgres
Username	The database account used to connect to the AnalyticDB for PostgreSQL instance. For more information about how to create an account, see Create a database account .	testuser
Password	The password of the database account used to connect to the AnalyticDB for PostgreSQL instance.	PassW0rd

5. Click **Test Connection**. A **Connected** message indicates that the connection is established.

Note If an error message of `org.postgresql.Driver` is returned, you can perform the following operations to **download or update** the driver library of Greenplum:

- On the **Main** tab of the **Connection Settings** page, click **Edit Driver Settings**.
- On the **Libraries** tab, click **Download/Update**.

6. Click **Finish**.

JDBC

You can use the PostgreSQL or Greenplum JDBC driver to connect to AnalyticDB for PostgreSQL.

Use one of the following methods to download the JDBC driver:

- Click [PostgreSQL JDBC Driver](#) to download the PostgreSQL JDBC driver and add it to an environment variable.
- Obtain the Greenplum JDBC driver from the Greenplum official website. For more information, see [Greenplum Database 4.3 Connectivity Tools for UNIX](#) or [Greenplum Client and Loader Tools Package](#).

Sample code:

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.SQLException;
import java.sql.Statement;
public class gp_conn {
    public static void main(String[] args) {
        try {
            Class.forName("org.postgresql.Driver");
            Connection db = DriverManager.getConnection("jdbc:postgresql://mygpdbpub.gpdb.rds.aliyuncs.com:5432/postgres", "mygpdb", "mygpdb");
            Statement st = db.createStatement();
            ResultSet rs = st.executeQuery("select * from gp_segment_configuration;");
            while (rs.next()) {
                System.out.print(rs.getString(1));
                System.out.print(" | ");
                System.out.print(rs.getString(2));
                System.out.print(" | ");
                System.out.print(rs.getString(3));
                System.out.print(" | ");
                System.out.print(rs.getString(4));
                System.out.print(" | ");
                System.out.print(rs.getString(5));
                System.out.print(" | ");
                System.out.print(rs.getString(6));
                System.out.print(" | ");
                System.out.print(rs.getString(7));
                System.out.print(" | ");
                System.out.print(rs.getString(8));
                System.out.print(" | ");
                System.out.print(rs.getString(9));
                System.out.print(" | ");
                System.out.print(rs.getString(10));
                System.out.print(" | ");
                System.out.println(rs.getString(11));
            }
            rs.close();
            st.close();
        } catch (ClassNotFoundException e) {
            e.printStackTrace();
        } catch (SQLException e) {
            e.printStackTrace();
        }
    }
}
```

For more information about the JDBC driver, see [The PostgreSQL JDBC Interface](#).

Python

You can use `psycopg2` to connect to Greenplum or PostgreSQL in Python. To use `psycopg2`, perform the following steps:

1. Install psycopg2. If you use CentOS, you can use one of the following installation methods:

- o Method 1: Run the `yum -y install python-psycopg2` command.
- o Method 2: Run the `pip install psycopg2` command.
- o Method 3: Use the source code.

```
yum install -y postgresql-devel*
wget http://initd.org/psycopg/tarballs/PSYCOPG-2-6/psycopg2-2.6.tar.gz
tar xf psycopg2-2.6.tar.gz
cd psycopg2-2.6
python setup.py build
sudo python setup.py install
```

2. Set the PYTHONPATH environment variable for subsequent variable reference. Sample code:

```
import psycopg2
sql = 'select * from gp_segment_configuration;'
conn = psycopg2.connect(database='gpdb', user='mygpdb', password='mygpdb', host='mygpdbpub.gpdb.rds.aliyuncs.com', port=5432)
conn.autocommit = True
cursor = conn.cursor()
cursor.execute(sql)
rows = cursor.fetchall()
for row in rows:
    print row
conn.commit()
conn.close()
```

Sample command output:

```
(1, -1, 'p', 'p', 's', 'u', 3022, '192.168.2.158', '192.168.2.158', None, None)
(6, -1, 'm', 'm', 's', 'u', 3019, '192.168.2.47', '192.168.2.47', None, None)
(2, 0, 'p', 'p', 's', 'u', 3025, '192.168.2.148', '192.168.2.148', 3525, None)
(4, 0, 'm', 'm', 's', 'u', 3024, '192.168.2.158', '192.168.2.158', 3524, None)
(3, 1, 'p', 'p', 's', 'u', 3023, '192.168.2.158', '192.168.2.158', 3523, None)
(5, 1, 'm', 'm', 's', 'u', 3026, '192.168.2.148', '192.168.2.148', 3526, None)
```

libpq

libpq is the C language interface to PostgreSQL. You can use the libpq library to connect to and manage PostgreSQL databases in a C program. If Greenplum or PostgreSQL is deployed, you can find both the static and dynamic libraries of libpq in the *lib* directory.

- For more information about examples of libpq, see [Example Programs](#).
- For more information about libpq, see [libpq - C Library](#).

ODBC

The PostgreSQL ODBC driver is an open source tool licensed based on the GNU Lesser General Public License (LGPL) protocol. You can download the driver from the [PostgreSQL official website](#).

To use the ODBC driver to connect to an AnalyticDB for PostgreSQL instance, perform the following steps:

1. Run the following commands to install the ODBC driver:


```
yum install -y unixODBC.x86_64
yum install -y postgresql-odbc.x86_64
```

2. Run the following command to configure the data source:

```
vim /etc/odbc.ini
```

Add connection information. Example:

```
[mygpdb]
Description = Test to gp
Driver = PostgreSQL
Database = ****
Servername = ****.gpdb.rds.aliyuncs.com
UserName = ****
Password = ****
Port = ****
ReadOnly = 0
```

 **Note** Replace `****` in the sample code with the corresponding connection information.

3. Test the connectivity. Example:

```
echo "select count(*) from pg_class" | isql mygpdb
+-----+
| Connected! |
|          |
| sql-statement |
| help [tablename] |
| quit |
|          |
+-----+
SQL> select count(*) from pg_class
+-----+
| count |
+-----+
| 388 |
+-----+
SQLRowCount returns 1
1 rows fetched
```

4. After the ODBC driver is connected to an instance, connect your application to the driver. For more information, see [PostgreSQL ODBC driver](#) and [psqlODBC HOWTO - C#](#).

Other client tools

- GUI client tools

You can use other GUI client tools supported by Greenplum, such as [DBeaver](#), to connect to an AnalyticDB for PostgreSQL instance.

- Greenplum client tools

The Greenplum official website provides a tool package that includes JDBC, ODBC, and libpq. This package is easy to install and use. For more information, see the [Greenplum documentation](#).

References

- [Pivotal Greenplum Database 4.3.33 Documentation](#)
- [Pivotal Greenplum 6.3 Documentation](#)
- [PostgreSQL ODBC driver](#)
- [Compiling psqLODBC on UNIX](#)
- [Greenplum ODBC Driver](#)
- [Greenplum JDBC Driver](#)