

CloudTable Service

Getting Started

Issue 01
Date 2024-04-30



Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: <https://www.huawei.com>

Email: support@huawei.com

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Getting Started with HBase.....	1
2 Getting Started with ClickHouse.....	7

1 Getting Started with HBase

HBase is a column-based distributed storage system that features high reliability, performance, and scalability. This section describes how to use HBase from scratch. For example, how to run the HBase shell command to create tables, insert data into tables, modify tables, read and delete table data, and delete tables.

Background Information

Suppose a user develops an application to manage users who use service A in an enterprise. The procedure of operating service A on the HBase client is as follows:

- Create the **user_info** table.
- Add users' educational backgrounds and professional titles to the table.
- Query user names and addresses by user ID.
- Query information by user name.
- Deregister users and delete user data from the user information table.
- Delete the user information table after service A ends.

Table 1-1 User information

ID	Name	Gender	Age	Address
12005000201	A	Male	19	Shenzhen, Guangdong
12005000202	B	Female	23	Shijiazhuang, Hebei
12005000203	C	Male	26	Ningbo, Zhejiang
12005000204	D	Male	18	Xiangyang, Hubei
12005000205	E	Female	21	Shangrao, Jiangxi
12005000206	F	Male	32	Zhuzhou, Hunan
12005000207	G	Female	29	Nanyang, Henan
12005000208	H	Female	30	Kaixian, Chongqing
12005000209	I	Male	26	Weinan, Shaanxi

ID	Name	Gender	Age	Address
12005000210	J	Male	25	Dalian, Liaoning

Buying an HBase Cluster

1. Log in to the CloudTable console.
2. Select a region in the upper left corner.
3. Click **Cluster Management**.
4. Click **Buy Cluster** in the upper right corner of the **Cluster Management** page and set related parameters. For details about how to configure ports for security group rules, see [Table 1-2](#).
5. Click **Buy Now**. On the displayed page, confirm the specifications and click **Finish**.
6. Return to the cluster list to view the cluster creation progress. If the cluster status is **In service**, the cluster is created. For details, see [Creating an HBase Cluster](#).

Table 1-2 Custom security rules

Direction	Protocol	Port/Range	Source/Security Group	Usage
Outbound	All	All	0.0.0.0/0	Permit in the outbound direction
Inbound	TCP	16000	Security group of the CloudTable HBase cluster	HMaster RPC port
	TCP	16020		RegionServer RPC port
	TCP	2181		ZooKeeper client connection monitoring port
	TCP	2888		Follower connection monitoring port
	TCP	3888		ZooKeeper election port
	TCP	2000		HAgent access port

Downloading the HBase Client

1. Log in to the CloudTable console.
2. Select a region in the upper left corner.
3. Click **Help** in the navigation pane.
4. Choose **Download the HBase x.x Client** under **Helpful Links** (choose the version that you need) on the right of the help page to download the client installation package.

Preparing an ECS


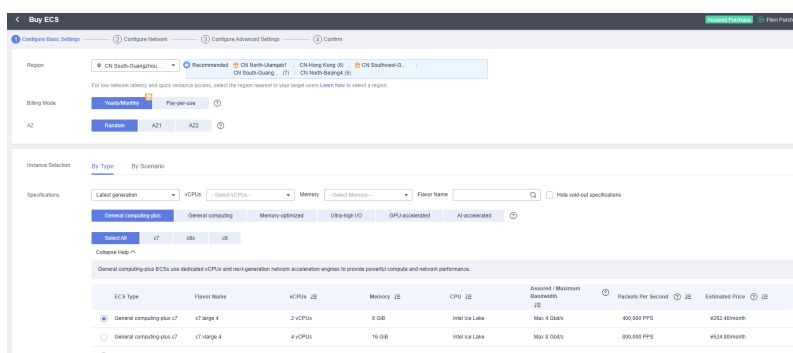

1. Buy an ECS.
 - a. Log in to the ECS console.
 - b. Click  in the upper left corner to select a region.
 - c. Click the service list icon and choose **Compute > Elastic Cloud Server**.
 - d. Click **Buy ECS**.
 - e. Configure ECS parameters.
 - i. Select the CPU architecture and specifications.

Figure 1-1 Configuring basic settings



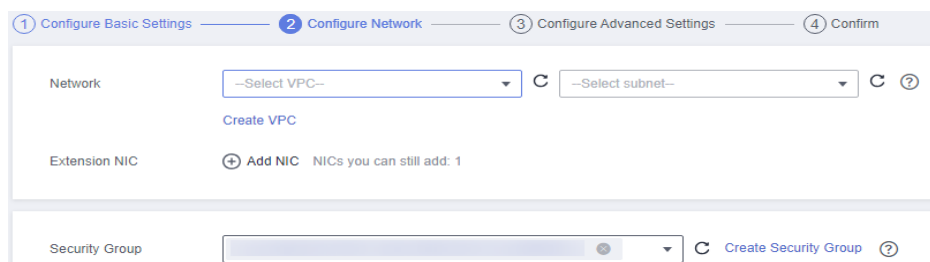
- ii. Select the image and disk specifications.

Figure 1-2 Selecting the image and disk specifications



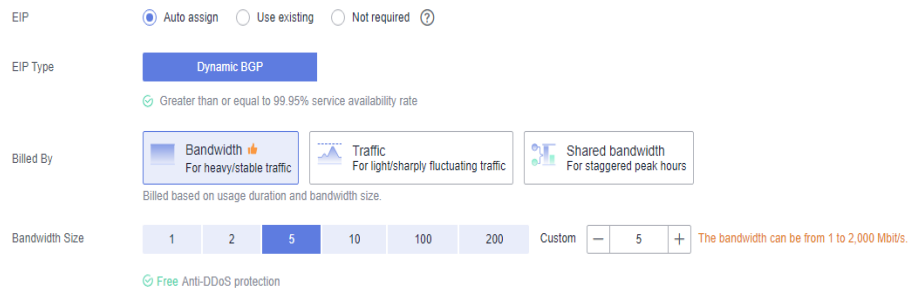
2. Click **Next: Configure Network**.
 - a. Select a VPC and a security group.

Figure 1-3 Configuring network



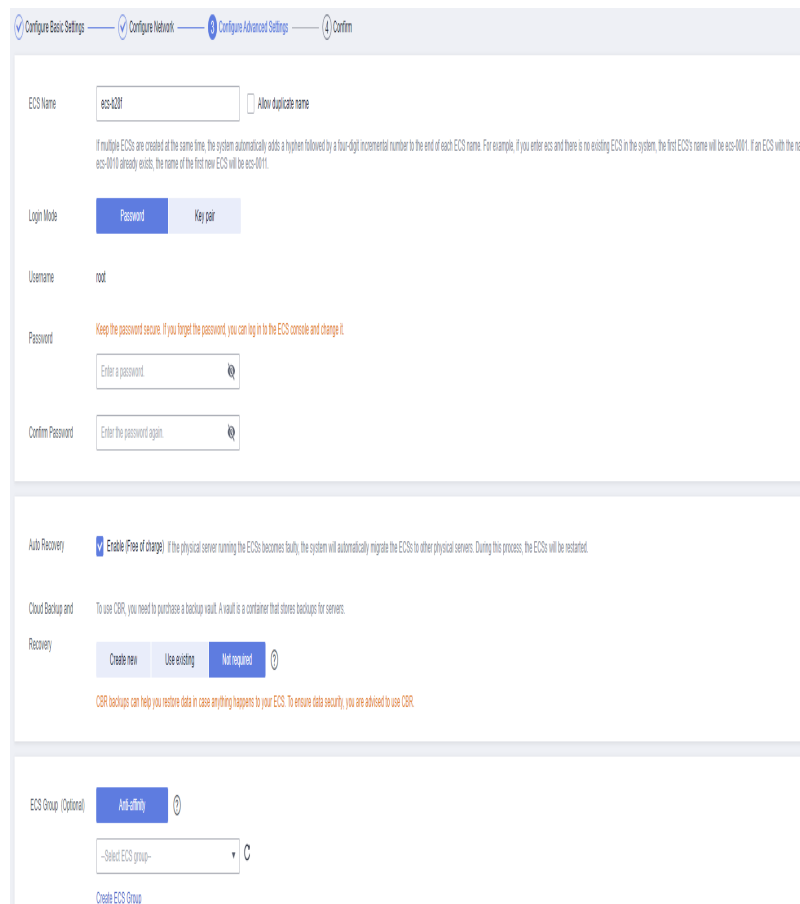
- b. Select **Auto assign** and set **Billed By** to **Traffic**.

Figure 1-4 Configuring EIP type and bandwidth billing type



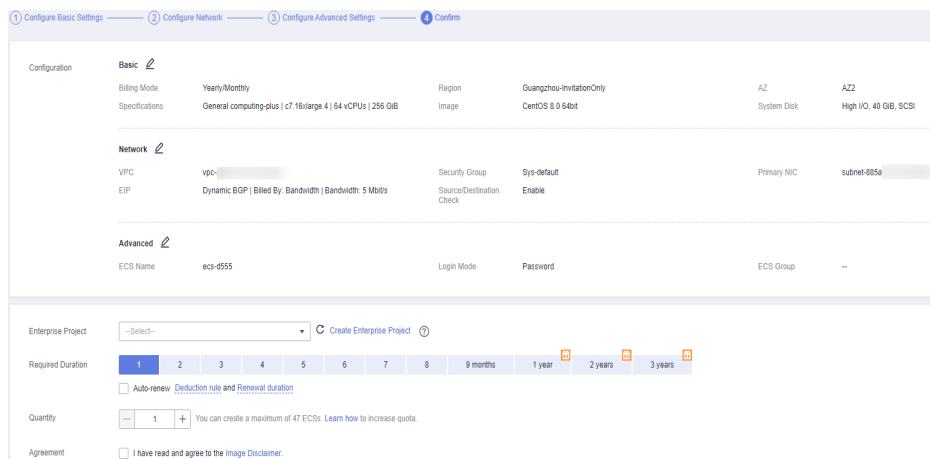
3. Click **Next: Configure Advanced Settings**.
Configure the ECS name and password.

Figure 1-5 Configuring advanced settings



4. Click **Next: Confirm** to confirm the configuration information.

Figure 1-6 Confirming the configuration information



5. Select an enterprise project, select the sentence in the **Agreement** area, confirm the configuration information, and click **Submit**.
6. Return to the ECS list page and view the ECS creation progress. When the ECS status changes to **Running**, the ECS is created.

Installing the Client

This part introduces how to manually install the client. You can also choose the one-click client deployment method.

1. Use the SSH login tool (such as PuTTY) to log in to the Linux ECS through the EIP.

For details about how to log in to the ECS, see "Remotely Logging In to a Linux ECS (Using an SSH Password)" in [Logging In to a Linux ECS](#) of the *Elastic Cloud Server User Guide*.

2. Transfer the client downloaded in 2 to the Linux ECS.
3. Decompress the client installation package.

```
cd <Path of the client installation package>
tar xzvf hbase-1.3.1-bin.tar.gz
cd <Path of the client installation package>
tar xzvf hbase-2.4.14-bin.tar.gz
```

<Path of the client installation package>. Replace it with the actual path.

4. Configure the ZooKeeper address in a configuration file.

In the decompression directory of the client installation package, open the **hbase/conf/hbase-site.xml** file and set the following parameters:

- a. **hbase.zookeeper.quorum**: The value of this parameter is the ZooKeeper link of the cluster found in the cluster list.

Log in to the CloudTable console and choose **Cluster Management**. In the cluster list, locate the required cluster and obtain its ZooKeeper link in the **ZK Link** column. Add this link address to the **hbase/conf/hbase-site.xml** file, as shown in [Figure 1-8](#).

Figure 1-7 Viewing the ZooKeeper link




Figure 1-8 Configuring the ZooKeeper address

```
<configuration>
<property>
<name>hbase.zookeeper.quorum</name>
<value>xxx-zk1.cloudtable.com,xxx-zk2.cloudtable.com,xxx-zk3.cloudtable.com</value>
</property>
```

- b. **mapreduce.cluster.local.dir**: Check whether the configuration item exists. If the configuration item does not exist, add it.

The configuration example is as follows:

```
<configuration>
<property>
<name>hbase.zookeeper.quorum</name>
<value>xxx-zk1.cloudtable.com,xxx-zk2.cloudtable.com,xxx-zk3.cloudtable.com</value>
</property>

<property>
<name>mapreduce.cluster.local.dir</name>
<value>${hadoop.tmp.dir}/mapred/local</value>
</property>
</configuration>
```

5. Start the shell to access the cluster. Run the **bin/hbase shell** command to start the shell to access the cluster.

Running the HBase Client Command to Implement Service A

1. Create the **user_info** table according to [Table 1-1](#) and add related data to it.

```
create 'user_info',{NAME => 'i'}
```

For example, to add information about the user whose ID is **12005000201**, run the following commands:

```
put 'user_info','12005000201','i:name','A'
put 'user_info','12005000201','i:gender','Male'
put 'user_info','12005000201','i:age','19'
put 'user_info','12005000201','i:address','Shenzhen, Guangdong'
```

2. Add users' educational backgrounds and titles to the **user_info** table.

For example, to add educational background and title information about user **12005000201**, run the following commands:

```
put 'user_info','12005000201','i:degree','master'
put 'user_info','12005000201','i:pose','manager'
```

3. Query user names and addresses by user ID.

For example, to query the name and address of user **12005000201**, run the following command:

```
scan 'user_info',
{STARTROW=>'12005000201',STOPROW=>'12005000201',COLUMNS=>['i:name','i:address']}
```

4. Query information by user name.

For example, to query information about user A, run the following command:

```
scan 'user_info',{FILTER=>"SingleColumnValueFilter('i','name',=,'binary:A')"
```

5. Delete user data from the user information table.

All user data needs to be deleted. For example, to delete data of user **12005000201**, run the following command:

```
delete 'user_info','12005000201','i'
```

6. Delete the user information table.

```
disable 'user_info';drop 'user_info'
```

2 Getting Started with ClickHouse

ClickHouse offers easy-to-use, flexible, and stable hosting services in the cloud. A data warehouse can be created in minutes for massive real-time data query and analysis, improving the overall efficiency of data value mining. By leveraging the massively parallel processing (MPP) architecture, ClickHouse can query data several times faster than conventional data warehouses.

Background Information

Assume that there is a student score table and you need to use ClickHouse to perform the following operations:

- Create the user information table **demo_t**.
- Add the user gender and subject to the user information.
- Query basic user information by user ID.
- Delete the user information table after the service ends.

Table 2-1 Score table

user_id	name	sex	subject	score	time
10000	A	1	Spanish	89	2023-07-01 09:00:00
10001	B	0	Math	132	2023-07-01 09:00:00
10002	C	0	Math	90	2023-07-02 09:00:00
10003	D	0	English	120	2023-07-01 14:00:00
10004	E	1	Spanish	101	2023-07-01 09:00:00
10005	F	1	Spanish	110	2023-07-01 09:00:00

Table 2-2 Parameters

Parameter	Description
10000	User ID, which uniquely identifies a user.
2023-07-01 09:00:00	Data import time.
A	Student name.
1	Gender female (0 indicates male).
Spanish	Discipline.
89	Score.

Buying a ClickHouse Cluster

1. Log in to the CloudTable console.
2. Select a region in the upper left corner.
3. Click **Cluster Management**.
4. Click **Buy Cluster** in the upper right corner of the **Cluster Management** page and set related parameters.
5. Click **Buy Now**. On the displayed page, confirm the specifications and click **Finish**.
6. Return to the cluster list to view the cluster creation progress. If the cluster status is **In service**, the cluster is created. For details, see [Creating a ClickHouse Cluster](#).

Downloading the ClickHouse Client

1. Log in to the CloudTable console.
2. Select a region in the upper left corner.
3. Click **Help** in the navigation pane.
4. Choose **Download the ClickHouse Client** under **Helpful Links** on the right of the help page to download the client installation package.

Preparing an ECS


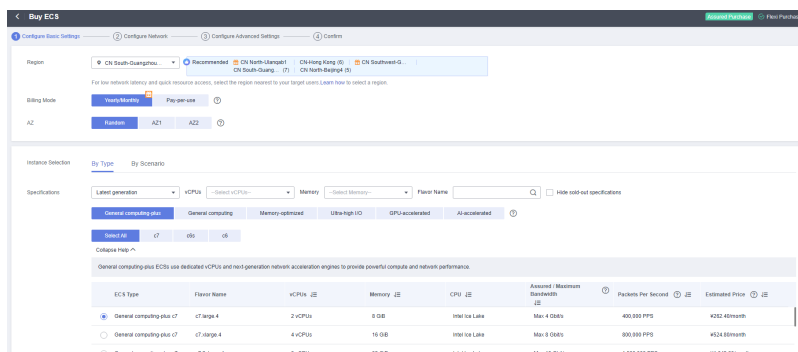

1. Buy an ECS.
 - a. Log in to the ECS console.
 - b. Click  in the upper left corner to select a region.
 - c. Click the service list icon and choose **Compute > Elastic Cloud Server**.
 - d. Click **Buy ECS**.
 - e. Configure ECS parameters.
 - i. Select the CPU architecture and specifications.

Figure 2-1 Configuring basic settings



ii. Select the image and disk specifications.

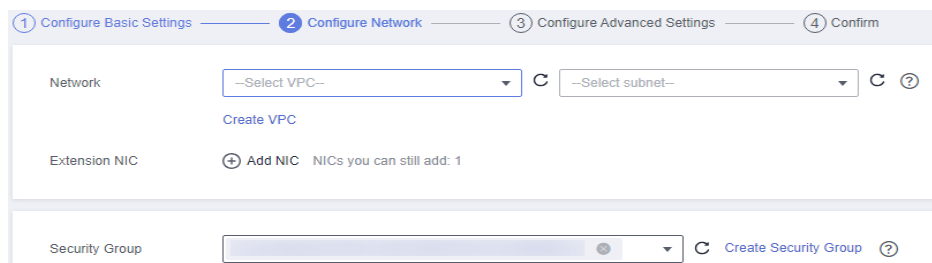
Figure 2-2 Selecting the image and disk specifications



2. Click **Next: Configure Network**.

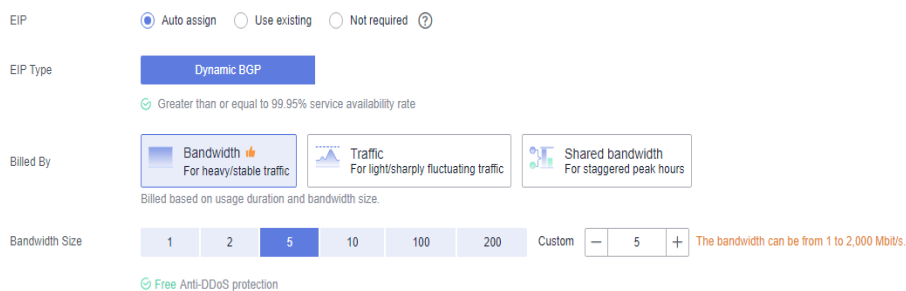
a. Select a VPC and a security group.

Figure 2-3 Configuring network



b. Select **Auto assign** and set **Billed By** to **Traffic**.

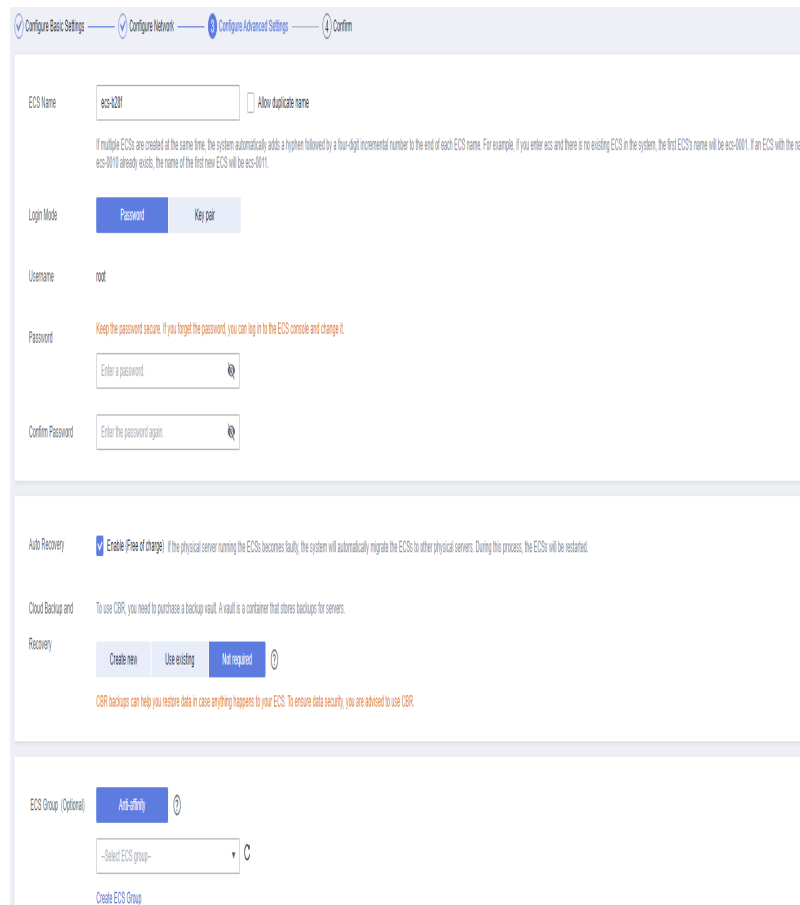
Figure 2-4 Configuring EIP type and bandwidth billing type



3. Click **Next: Configure Advanced Settings**.

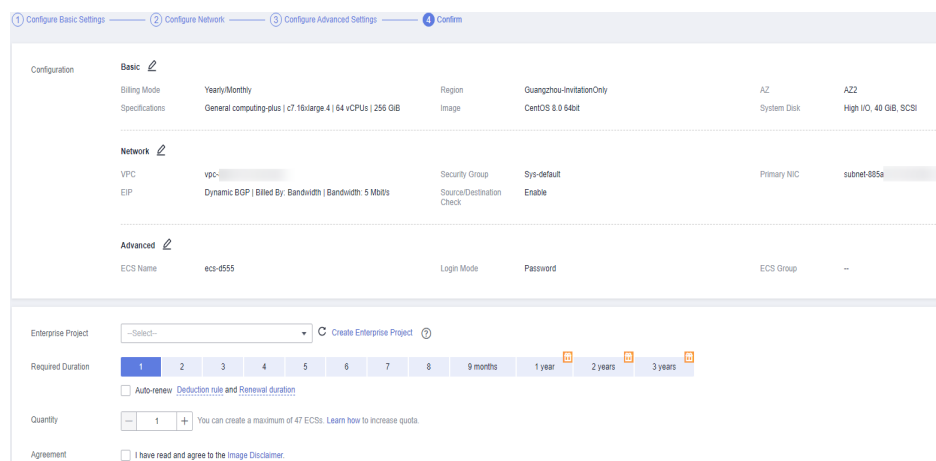
Configure the ECS name and password.

Figure 2-5 Configuring advanced settings



4. Click **Next: Confirm** to confirm the configuration information.

Figure 2-6 Confirming the configuration information



5. Select an enterprise project, select the sentence in the **Agreement** area, confirm the configuration information, and click **Submit**.
6. Return to the ECS list page and view the ECS creation progress. When the ECS status changes to **Running**, the ECS is created.

Installing the ClickHouse Client

You can manually install the client on an ECS.

1. Use the SSH login tool (such as PuTTY) to log in to the Linux ECS through the EIP.

For details about how to log in to the ECS, see "Remotely Logging In to a Linux ECS (Using an SSH Password)" in [Logging In to a Linux ECS](#) of the *Elastic Cloud Server User Guide*.

2. Transfer the client downloaded in 2 to the Linux ECS.
3. Install the client and connect to the cluster.

- a. Use the SSH login tool to remotely log in to the Linux ECS through the EIP.

For details, see [Login Using an SSH Password](#) in the *Elastic Cloud Server User Guide*.

- b. Go to the root directory of the SSH login tool.

```
cd /
```

- c. Create a folder in the root directory.

```
mkdir Folder name
```

- d. Go to the directory of the created folder.

```
cd /Folder name/
```

- e. Place the client in the directory.

- f. Decompress the client package.

```
tar -zxf Client package name
```

- g. Load the **.so** file.

```
sh install.sh
```

- h. Go to the **bin** directory.

```
cd bin/
```

Grant the 700 permission to the directory.

```
chmod 700 clickhouse
```

- i. Connect to the port of the ClickHouse cluster.

Use the following command to connect to a normal cluster.

```
./clickhouse client --host Private IP address of the cluster --port Port --user admin --password Password
```

Check the security cluster connection commands in [Connecting to a Security Cluster](#).

```
./clickhouse client --host Private IP address of the cluster --port port --user admin --password Password --secure --config-file /root/config.xml
```

NOTE

- *Private IP address of the cluster*: cluster access address on the cluster details page. Replace it with the access address of the cluster you purchased.
- *Password*: the password set when you purchase the cluster. If there are special characters, use backslashes (\) to escape them. If the password is enclosed in single quotation marks ('), the special characters do not need to be escaped.

Table 2-3 Custom security rules

Direction	Action	Port/Range	Type	Destination/Source Address	Usage	
Outbound	Allow	All	IPv4/IPv6	0.0.0.0/0	Permit in the outbound direction	
Inbound	Allow	8123		Security group of the CloudTable ClickHouse cluster		ClickHouse HTTP port number
	Allow	9000				ClickHouse TCP port number
	Allow	8443				ClickHouse HTTPS port number
	Allow	9440				Secure TCP security port of ClickHouse

Inserting Data

Create a data table using the ClickHouse cluster and query the table data.

1. Create a database.

```
create database DB_demo;
```

2. Use the database.

```
use DB_demo;
```

3. Create a table.

```
create table DB_demo_t(user_id Int32,name String,sex Tinyint ,subject String,score Int32,time datetime)engine=TinyLog;
```

4. Insert data.

```
insert into DB_demo_t(user_id,name,sex,subject,score,time)
values('10000','A','1','Spanish','89','2023-07-01 09:00:00');
insert into DB_demo_t(user_id,name,sex,subject,score,time)
values('10001','B','0','Math','132','2023-07-01 09:00:00');
insert into DB_demo_t(user_id,name,sex,subject,score,time)
values('10002','C','0','Math','90','2023-07-02 09:00:00');
insert into DB_demo_t(user_id,name,sex,subject,score,time)
values('10003','D','0','English','120','2023-07-01 14:00:00');
insert into DB_demo_t(user_id,name,sex,subject,score,time)
values('10004','E','1','Spanish','101','2023-07-01 09:00:00');
insert into DB_demo_t(user_id,name,sex,subject,score,time)
values('10005','F','1','Spanish','110','2023-07-01 09:00:00');
```

5. Query data.

- Query the imported data.

```
host-172-16-13-95 :) select * from DB_demo_t;
```

```
SELECT *
```

```
FROM DB_demo_t
```

```
Query id: 4e119f77-0592-4131-bbe2-31f42bc069a1
```

user_id	name	sex	subject	score	time
10000	A	1	Spanish	89	2023-07-01 09:00:00
10001	B	0	Math	132	2023-07-01 09:00:00
10002	C	0	Math	90	2023-07-02 09:00:00
10003	D	0	English	120	2023-07-01 14:00:00
10004	E	1	Spanish	101	2023-07-01 09:00:00
10005	F	1	Spanish	110	2023-07-01 09:00:00


```
└─┘  
6 rows in set. Elapsed: 0.004 sec.
```

6. Delete data.

- Delete the table.

```
drop table DB_demo_t;
```

- Delete the database.

```
drop database DB_demo;
```