

# **Tencent Cloud TCHouse-D**

## **Getting Started**

### **Product Documentation**



## Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

## Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

## Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

# Contents

## Getting Started

Using Tencent Cloud TCHouse-D Through the Console

Using Tencent Cloud TCHouse-D Through a Client

# Getting Started

## Using Tencent Cloud TCHouse-D Through the Console

Last updated : 2024-06-27 10:39:42

This document describes how to use Tencent Cloud TCHouse-D through the console to create, manage, and execute SQL commands in the SQL workspace.

### Creating Cluster

1. On the Tencent Cloud TCHouse-D introduction page, click **Buy Now**. Alternatively, log in to the Tencent Cloud TCHouse-D console and click **Create Cluster** to go to the purchase page for configuration and purchase.
2. After completing the basic configuration, cluster configuration, and log configuration, the system will display the corresponding configuration costs. Click **Activate** to create a cluster.

**Cluster**

Cluster name   
6-36 characters; supports Chinese characters, letters, digits, hyphens (-), and underscores (\_)

Kernel version

High availability (HA)  Non-High Availability  Read High Availability  Read/Write High Availability  
Deploy at least 3 FE nodes, including 1 node as a Follower node and the others as Observer nodes.

FE node type  Standard

Compute spec

Storage spec     
200 to 32000 GB for a single node

FE node count

The number of nodes cannot exceed the result of the number of available subnet IP addresses in the selected network minus the number of BE nodes. If the number of subnet IP addresses is not enough, switch the subnet or VPC and try again.

BE node type  Standard  High-performance

Compute spec

Storage spec     
200 to 32000 GB for a single node

BE node count

The number of nodes must be larger than or equal to 3 but cannot exceed the result of the number of available subnet IP addresses in the selected network minus the number of FE nodes. If the number of subnet IP addresses is not enough, switch the subnet or VPC and try again.

### View Cluster Information

1. After the cluster is created, you can access the [Tencent Cloud TCHouse-D console](#) and select the region where the cluster is located to view the status information of all clusters in that region from the cluster list.

**Cluster list** ap-hongkong 5 Other regions 5

[Create cluster](#) Separate search items with

Resource ID/Name	Status (progress)	Health status	FE node	BE node	Kernel version	AZ	Network	Subnet
cdwdoris-n3m4593c xcfxvsdf	Serving	Healthy	Standard, 3 node(s) 4-core 16GB, 200 GB	High IO, 3 node(s) 16-core 64 GB, 3570 GB	2.0	ap-hongkong-2	vpc-f0m0cm2m test	subnet-6n9oimbn 10.22.13.1

2. In the cluster list, click **Resource ID/Name** to view the basic cluster information, cluster status, configuration information, and network information.

**Cluster Info**

**Basic info**

Cluster ID: [ID]

Cluster name: [Name]

Billing mode: Pay-as-you-go

Creation time: 2024-04-08 17:40:44

Tag: [Change](#)

**Configuration info**

Kernel version: 2.0 (tencent-cdw-doris-2.0.7-f5a305b-0889e65)

High availability (HA): Read high availability

FE node spec: Standard, 4-core 16 GB / 3 node(s) / CLOUD\_HSSD 200 GB

BE node spec: High IO, 16-core 64 GB / 3 node(s) / LOCAL\_BASIC 3570 GB

Table names are case-sensitive: Case-Sensitive

Cloud Block Storage Encryption: Enable

Security Groups: None

**Cluster status**

Cluster status: Serving

**Network info**

AZ: ap-hongkong-2

VPC ID: [VPC ID]

Subnet ID: [Subnet ID]

JDBC access address: [Address]

HTTP connection address: [Address]

**Node info**

No.	Node type
1	FE(master)
2	FE(observer)
3	FE(observer)
4	BE
5	BE
6	BE

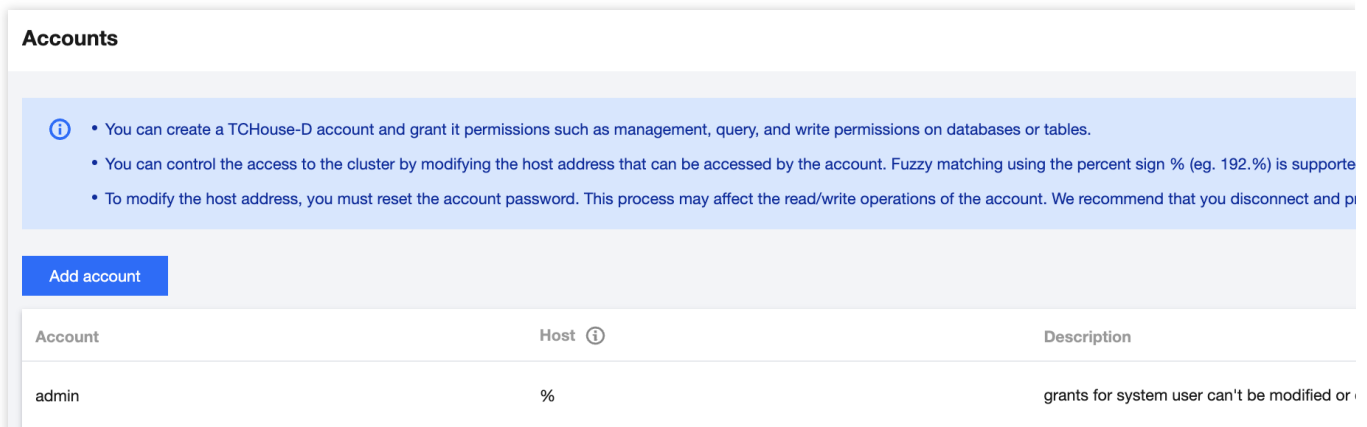
Total items: 6

## Cluster Operation

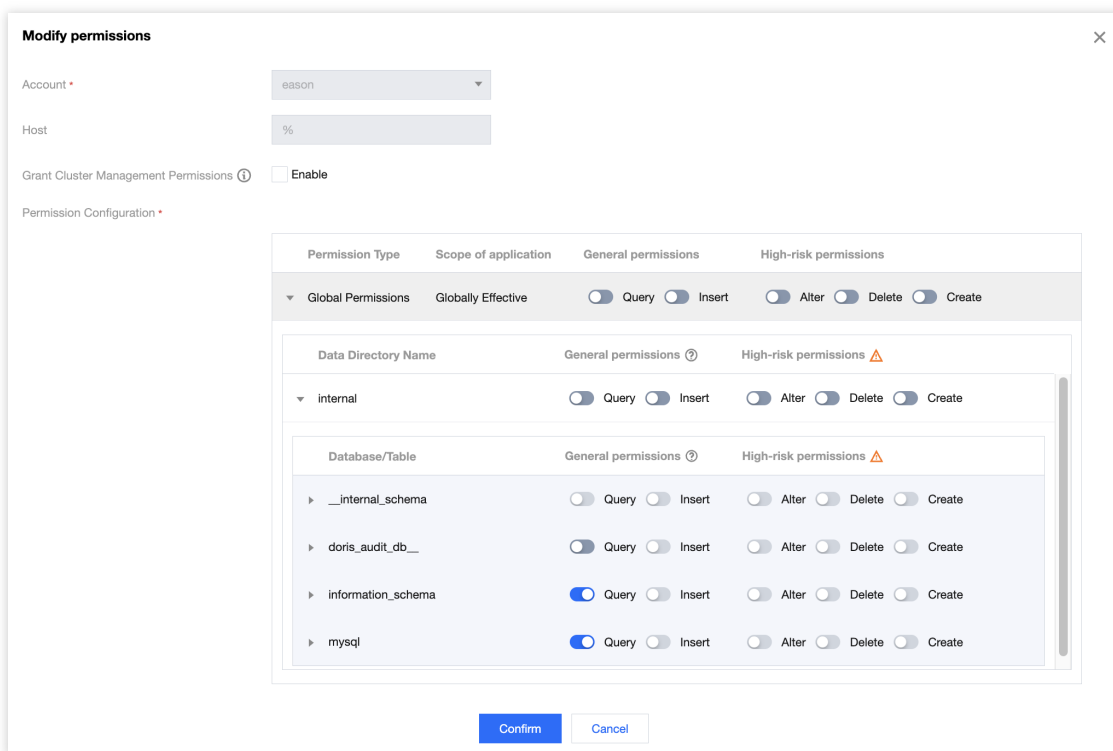
In the console, you can conveniently manage account permissions, configure monitoring and alarms, modify configurations, and manage nodes. For details, see the [Operation Guide](#) section.

## Managing Account Permissions

1. Tencent Cloud TCHouse-D offers a visual interface in the console, allowing for convenient and efficient account and permission management of clusters.

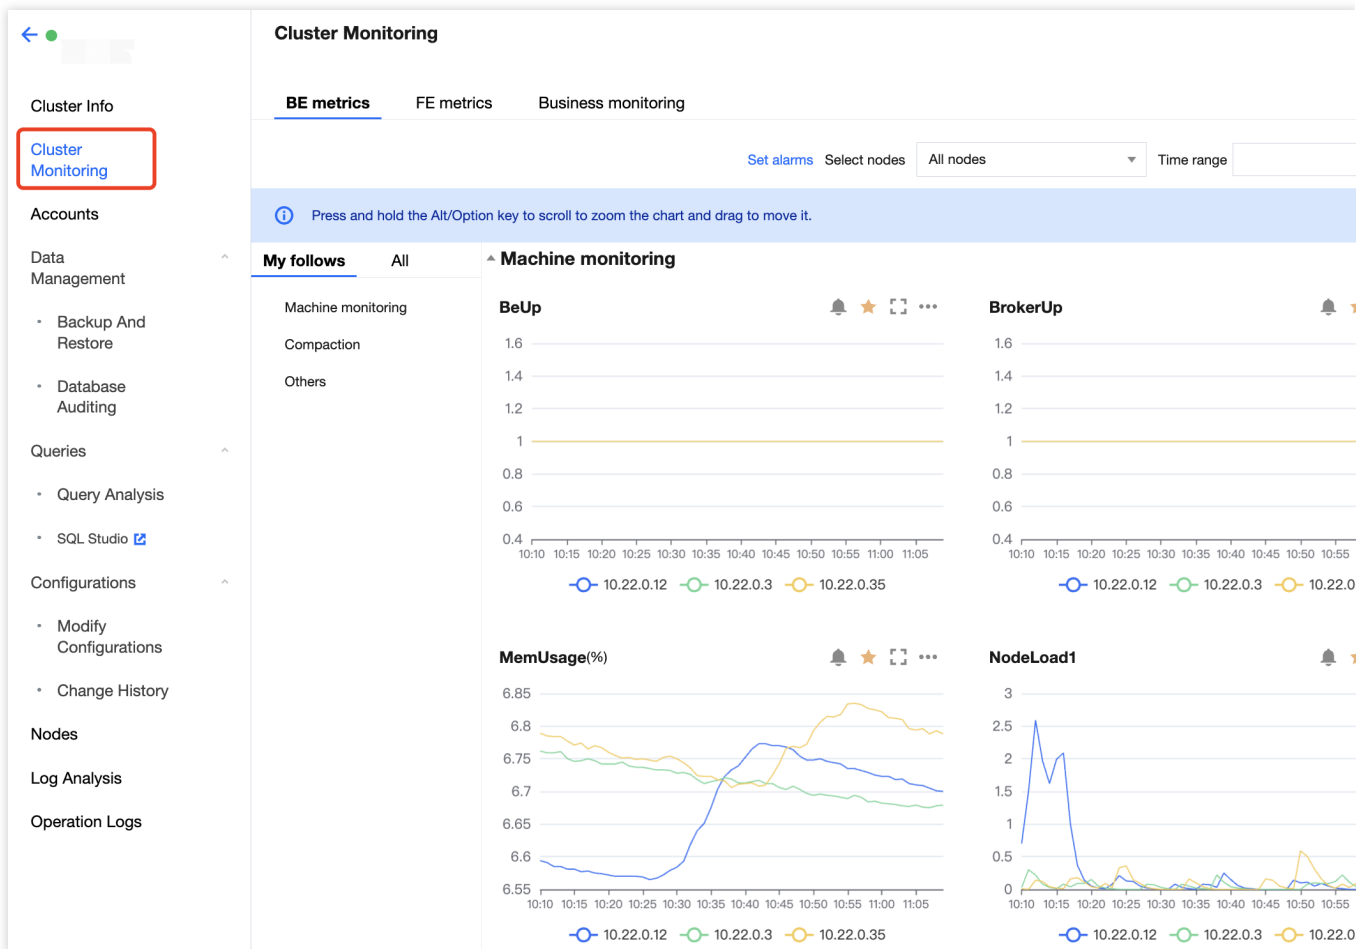


2. Click **Add Account**, fill in the database account, password, confirm password, and description (optional), and click **Confirm** to add an account. Click **Modify Permissions** to go to the permission modification window, where you can also view permissions.




### Monitoring and Alarm Configuration

1. Tencent Cloud TCHouse-D provides a wealth of monitoring indicators to help you understand the operation status of the cluster. You can configure alarms for indicators to check operation information in real time, achieving rapid response. Click **Resource ID/Name** of the cluster to view cluster details, and click **Cluster Monitoring** to view the performance indicators of the cluster.



2. Flexible and comprehensive alarm policies can be created, so that exception alarms can be reported for all monitoring indicators. In the cluster monitoring page, click **Set Alarm** to create a new alarm policy.



### Modifying Configuration

1. Tencent Cloud TCHouse-D supports flexible parameter configuration methods. You can comprehensively and intuitively set the parameters of BE/FE/Broker in the console, and view historical modification records. In the cluster list, click **Resource ID/Name** to view cluster details, and select **Configurations > Modify Configurations** from the left list to go to the parameter configuration page.

**Modify Configurations**

Configuration file Upload

Configuration file list:

- apache\_hdfs\_broker.conf
- be.conf
- fe.conf
- core-site.xml
- hdfs-site.xml
- hive-site.xml
- odbcinst.ini

Buttons: Apply to cluster, Undo all changes, Restart service

Parameter name	Reference value	Current Configuration Value	Value ran
broker_ipc_port ⓘ		8000	
client_expire_seconds ⓘ		300	
XMX ⓘ	2g	2g ✎	>=2g, <=4g

[Add parameter](#)

2. For the configuration parameters corresponding to Broker, BE, and FE, modifications to the configuration file and hot updates to the parameters are supported. An entire XML configuration file can be modified.

## Node Operations

1. The node management page displays the status list of each role in the cluster. The roles include FE, BE, and BROKER. The health status, process status, node IP address, and the most recent restart time (the last time the service was manually restarted on the role management page) are displayed for each role.



The screenshot shows the 'Nodes' management interface for a Doris cluster. The left sidebar is a navigation menu with 'Nodes' highlighted. The main content area shows a table of nodes with the following data:

Role	Node health status	Node IP
<input type="checkbox"/> FE(master)	Healthy	10.22.0.2
<input type="checkbox"/> BROKER	Healthy	10.22.0.2
<input type="checkbox"/> FE(observer)	Healthy	10.22.0.6
<input type="checkbox"/> BROKER	Healthy	10.22.0.6
<input type="checkbox"/> FE(observer)	Healthy	10.22.0.11
<input type="checkbox"/> BROKER	Healthy	10.22.0.11
<input type="checkbox"/> BE	Healthy	10.22.0.12
<input type="checkbox"/> BROKER	Healthy	10.22.0.12
<input type="checkbox"/> BE	Healthy	10.22.0.3
<input type="checkbox"/> BROKER	Healthy	10.22.0.3
<input type="checkbox"/> BE	Healthy	10.22.0.35
<input type="checkbox"/> BROKER	Healthy	10.22.0.35

2. You can restart each node. The restart methods include immediate restart, graceful restart, and rolling restart, and batch execution for the same type of nodes is supported.

## Executing SQL Commands

1. In the SQL workspace, you can quickly connect to the cluster and use SQL commands to perform a series of operations. For details, see [SQL Studio](#). Click on the console left list to access **SQL Studio**. In the cluster list, select the target **Resource ID/Name**. You can also find the entry under Query Management.

The screenshot displays the Tencent Cloud console interface for a TCHouse-D cluster. On the left, a dark sidebar contains navigation options: 'Overview', 'Clusters', and 'SQL Studio' (highlighted with a red box). The main content area is divided into two columns. The left column, under 'Cluster Info', lists various management options: 'Cluster Monitoring', 'Accounts', 'Data Management' (with a sub-item 'Backup And Restore'), 'Queries' (with sub-items 'Query Analysis' and 'SQL Studio' highlighted with a red box), 'Configurations' (with sub-items 'Modify Configurations' and 'Change History'), 'Nodes', 'Log Analysis', and 'Operation Logs'. The right column contains two panels: 'Basic info' and 'Configuration info'. The 'Basic info' panel lists: Cluster ID (with a copy icon), Cluster name 'xcfxvsdf' (with an edit icon), Billing mode 'Pay-as-you-go', Creation time '2024-04-08 17:40:44', and Tag 'Change' (with a link icon). The 'Configuration info' panel lists: Kernel version '2.0 (tencent-cdw-doris-2.0.7)', High availability (HA) 'Read high availability', FE node spec 'Standard, 4-core 16 GB / 3 nc', BE node spec 'High IO, 16-core 64 GB / 3 nc', Table names are case-sensitive 'Case-Sensitive', Cloud Block Storage Encryption 'Enable', and Security Groups 'None' (with an edit icon).

2. In the SQL studio, you can execute SQL statements, view execution records, and perform other operations.

The screenshot shows the Tencent Cloud console interface for a database instance. The top navigation bar includes the instance name 'xcfxvsdf - ap-hongkong'. The left sidebar shows 'Data sources' with 'internal' expanded, listing 'Information\_schema(default...)' and 'mysql(default\_cluster)'. The main area is divided into a 'Table' view on the left and an 'SQL 1' editor on the right. The 'Table' view displays a list of system tables such as 'character\_sets', 'collations', 'column\_privileges', etc. The 'SQL 1' editor shows a 'Running' status and a 'Format' button. Below the editor is an 'Execution log' section with a search filter and a table header. The table header includes columns for 'No.', 'Start time', 'Execution d...', 'Executi...', 'Node', 'Executing u...', and 'Executed SQL'. The table body is currently empty, showing 'No data yet' and 'Total items: 0'.

No.	Start time	Execution d...	Executi...	Node	Executing u...	Executed SQL
No data yet						

# Using Tencent Cloud TCHouse-D Through a Client

Last updated : 2024-06-27 10:40:06

Except the console, you can use Tencent Cloud TCHouse-D through a client based on JDBC and WebUI. After purchasing and creating a cluster, you can enter the cluster details page from [Cluster list](#) by clicking on **Resource ID/name**. The network information section provides two access addresses: JDBC and HTTP, which indicate the two types of clients:

The JDBC connection address is suitable for JDBC clients, such as commonly used MySQL clients.

The HTTP connection address is used to perform various operations on the cluster's REST API through the WebUI.

### Cluster Info

#### Basic info

Cluster ID

Cluster name  ✎

Billing mode Pay-as-you-go

Creation time 2024-04-08 17:40:44

Tag Change

#### Configuration info

Kernel version 2.0 (tencent-cdw-doris-2.0.7-f5a305b-0889e65)

High availability (HA) Read high availability

FE node spec Standard, 4-core 16 GB / 3 node(s) / CLOUD\_HSSD 200 GB

BE node spec High IO, 16-core 64 GB / 3 node(s) / LOCAL\_BASIC 3570 GB

Table names are case-sensitive. Case-Sensitive

Cloud Block Storage Encryption Enable

Security Groups None ✎

#### Cluster status

Cluster status Serving

#### Network info

AZ ap-hongkong-2

VPC ID

Subnet ID

JDBC access address

HTTP connection address

Node info

No.	Node type
1	FE(master)
2	FE(observer)
3	FE(observer)
4	BE
5	BE
6	BE

Total items: 6

## JDBC Clients (MySQL Client and More)

Tencent Cloud TCHouse-D uses the MySQL protocol for communication, and users can connect to the cluster through the MySQL Client or MySQL JDBC. It is recommended that you use the MySQL Client in versions later than 5.1 because the MySQL Client in versions prior to 5.1 does not support usernames longer than 16 characters.

### Note:

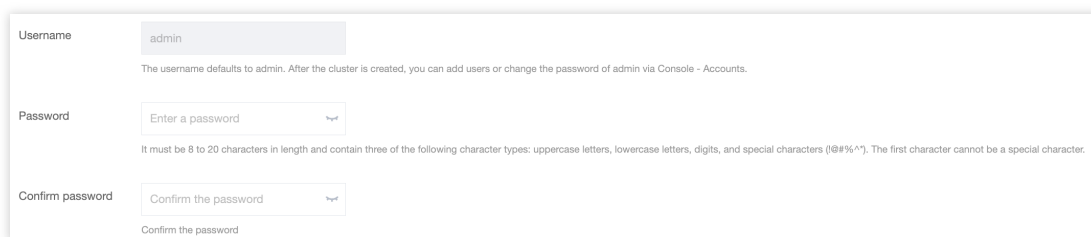
Any client that supports the MySQL JDBC protocol can be used, but there may be situations where the client does not display cluster's return information exactly as it is. Generally, this phenomenon does not cause problems, but operation results may be misjudged because some information is not displayed.

For example, Warning in the return information is not displayed in some clients after Broker Load is executed. Consequently, users may mistakenly think that there is no Warning (that is, no filtered data).

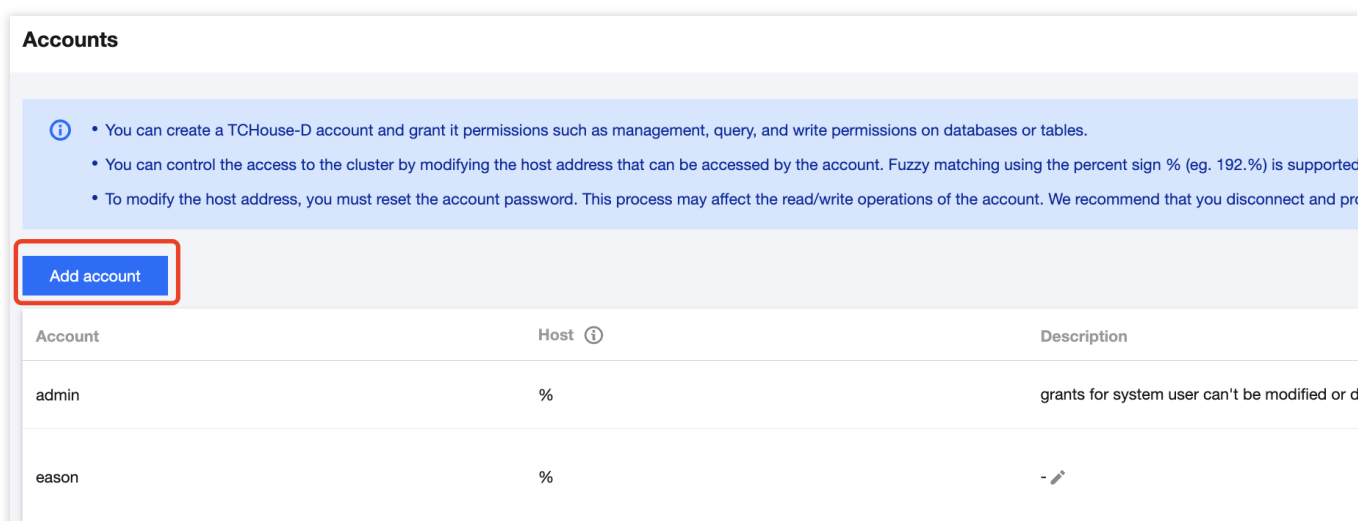
The following describes the basic process of connecting to Tencent Cloud TCHouse-D through MySQL Client.

### Step 1: Creating a User

1. When you purchase and create a Tencent Cloud TCHouse-D cluster, you need to set the admin user's password.



2. After the cluster is created, you can conveniently view cluster information and carry out operations such as creating new accounts in the console. You can use the admin user to connect to the cluster, or create a new user and use the sub-user to connect to the cluster.



Account	Host	Description
admin	%	grants for system user can't be modified or d
eason	%	-

3. When the cluster is running, you can use MySQL Client to connect to the cluster in a networked environment. In the cluster details page, the connection address can be copied with one click. If you want to use a sub-user to connect or

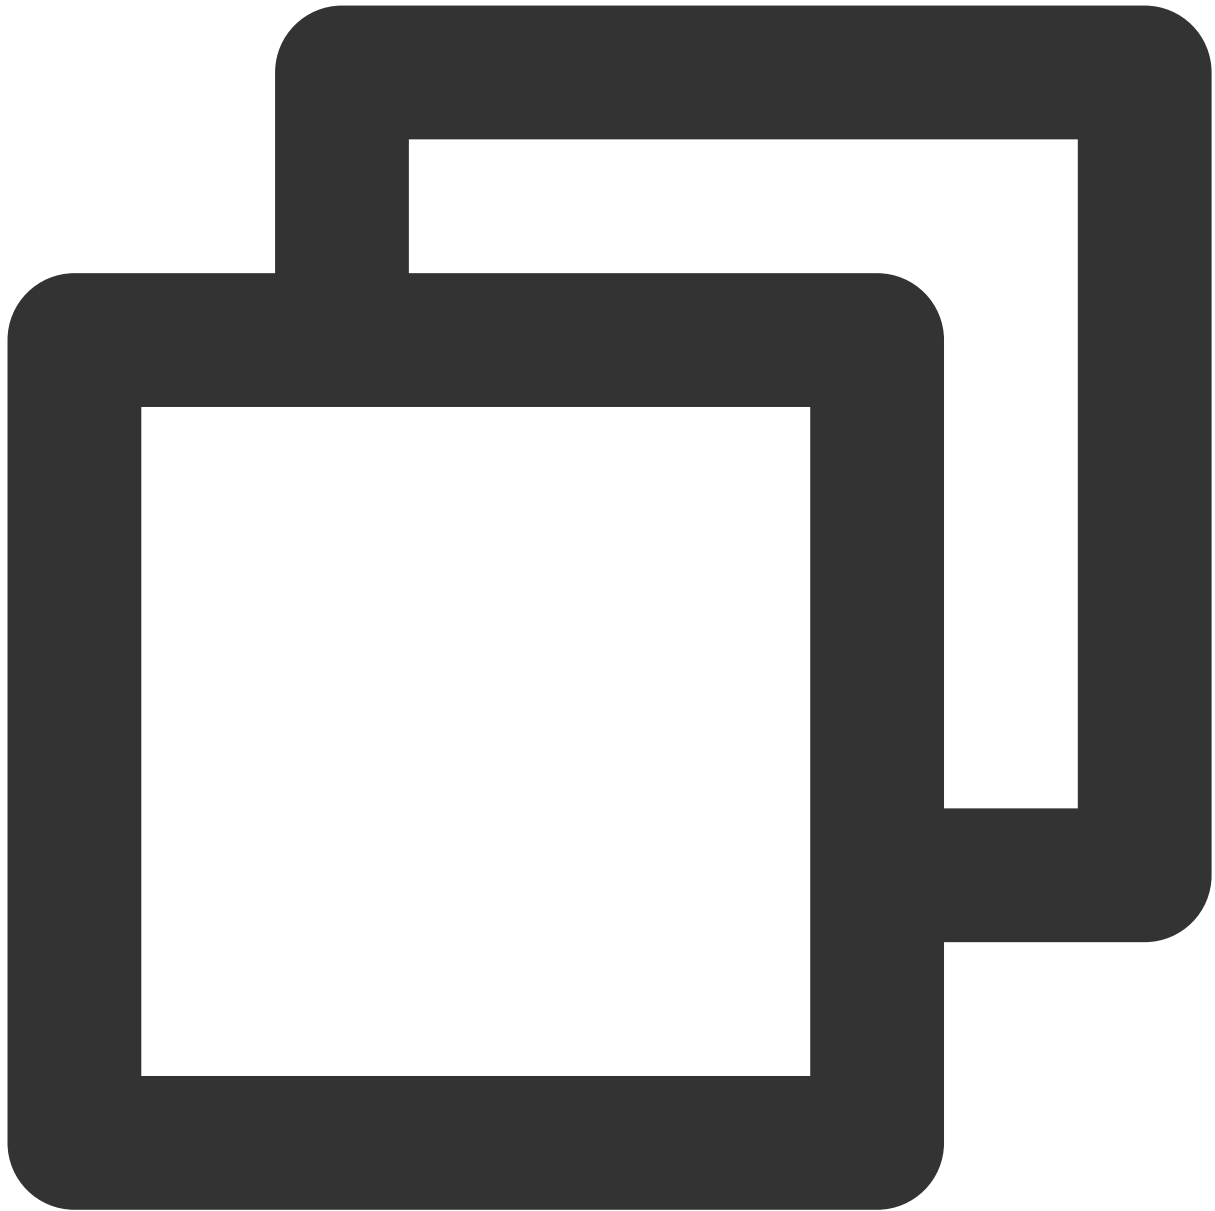
reset the password, you can do so on the **Accounts** page in the console.



**Network info**

AZ	ap-hongkong-2
VPC ID	mysql -h10.22.0.27 -P9030 -uadmin -p
Subnet ID	jdbc:mysql://10.22.0.27:9030
JDBC access address	10.22.0.27:9030 <a href="#">Copy connection string</a>
HTTP connection address	10.22.0.27:8030 <a href="#">Copy connection string</a>

4. Enter the following command to log in to the cluster:




```
mysql -h FE_HOST -P9030 -uadmin -p
```

**Note:**

FE\_HOST is the IP address of any FE node. 9030 is the query\_port configuration in fe.conf.

5. After logging in, you can also modify the admin password using the following command:



```
SET PASSWORD FOR 'admin' = PASSWORD('your_password');
```

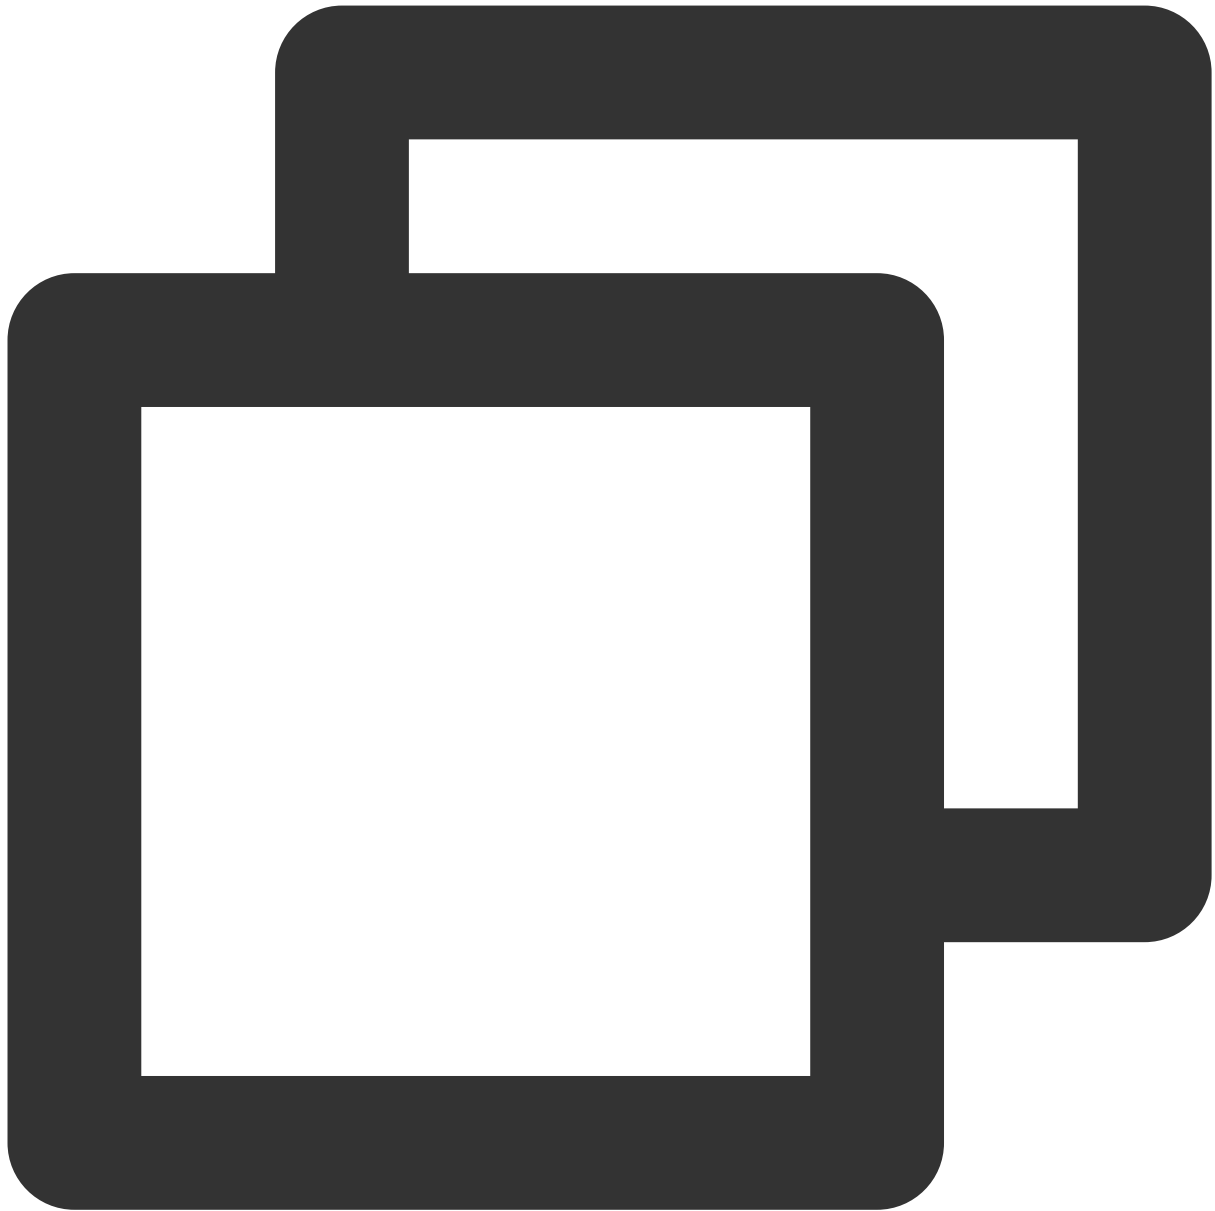
6. Create a new user:





```
CREATE USER 'test' IDENTIFIED BY 'test_passwd';
```

7. You can connect to the cluster through the newly created user using the following command:



```
mysql -h FE_HOST -P9030 -utest -ptest_passwd
```

**Note:**

By default, the newly created ordinary user has no permissions. You can grant permissions **Accounts** page in the console.

**Step 2: Creating a Data Table**

1. Create a database.


Initially, you can create a database through the admin user: `CREATE DATABASE example_db; .`

**Note:**

HELP command; can be used to check detailed syntax help for all commands. Example: `HELP CREATE DATABASE;` .

If you do not know the full name of a command, you can use help One field of the command for fuzzy queries. For example, if you type `HELP CREATE`, you can match commands such as `CREATE DATABASE` , `CREATE TABLE` , and `CREATE USER` .

After the database is created, you can check the database information using `SHOW DATABASES;` .



```
MySQL> SHOW DATABASES;
+-----+
```

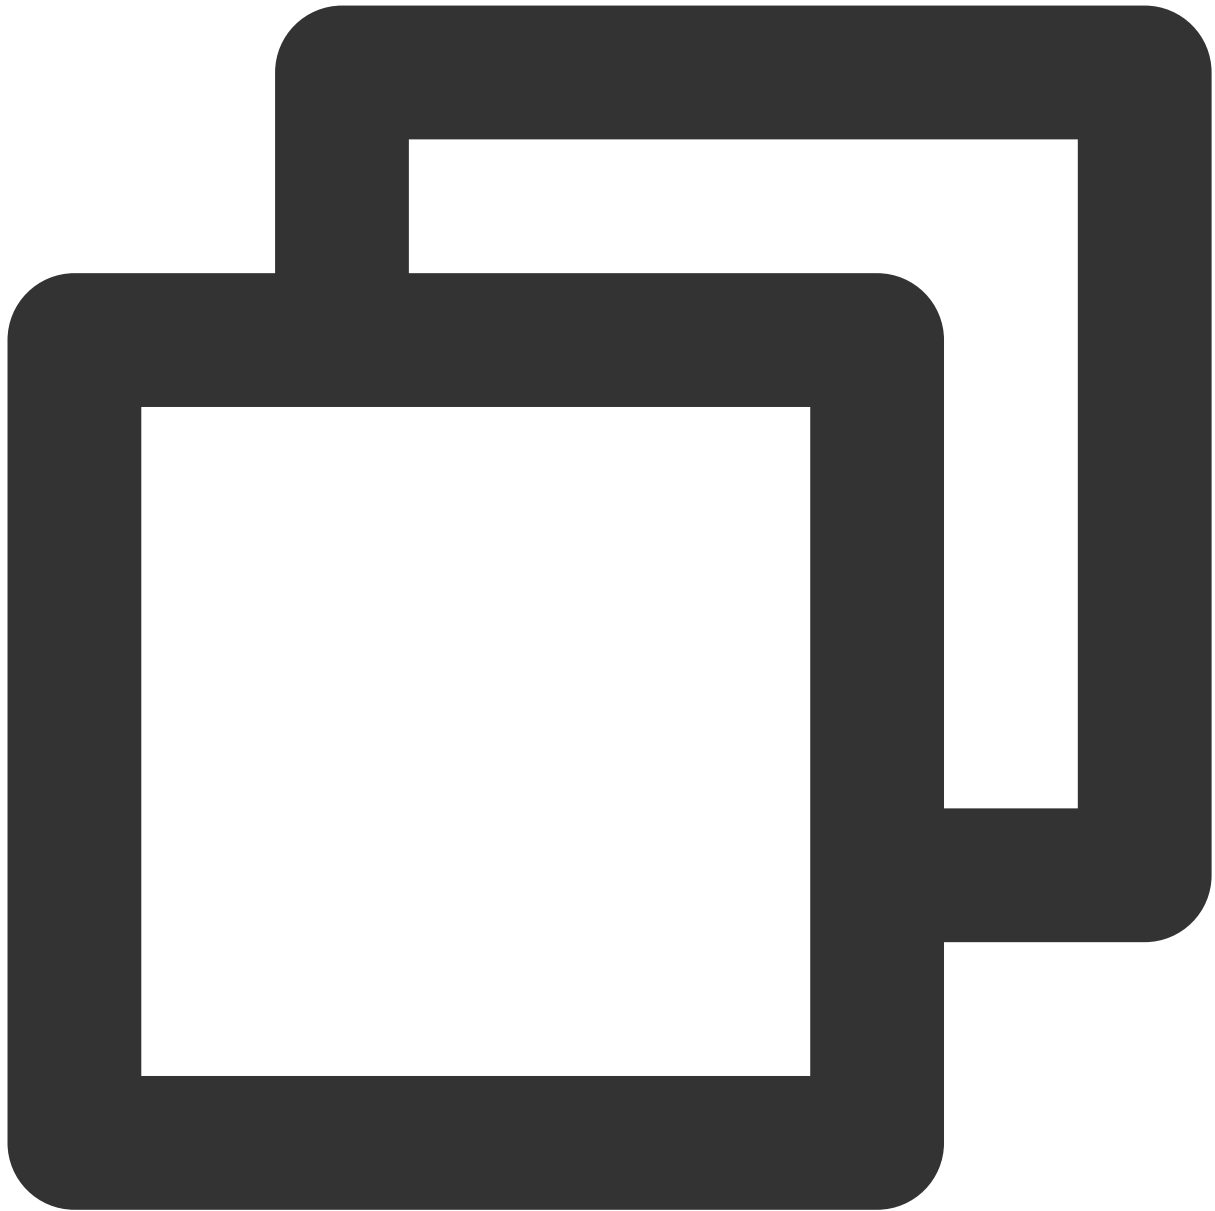
```
| Database          |
+-----+
| doris_audit_db__ |
| example_db       |
| information_schema |
+-----+
3 rows in set (0.00 sec)
```

information\_schema aims to be compatible with the MySQL protocol. Actual information may not be accurate.

Therefore, it is recommended that you obtain specific database information by querying the relevant database directly.

## 2. Grant the account permission.

After example\_db is created, you can use the admin account to grant the read/write permissions to the ordinary account, such as test. After the authorization, you can use the test account to operate on the example\_db database.



```
GRANT ALL ON example_db TO test;
```

### 3. Create a table.

First, switch the database:



```
USE example_db;
```

Use the `CREATE TABLE` command to create a table. Tencent Cloud TCHouse-D supports single partition and composite partition to create tables. For details, see [Data Partition and Bucketing](#). The following is an example of creating a table with aggregation model, showcasing both partition types.

### Single partition

Create a logical table named table1. The bucketing column is siteid, and the number of buckets is 10. The table's schema is as follows:


siteid: The type is INT (4 bytes), and the default value is 10.

citycode: The type is SMALLINT (2 bytes).

username: The type is VARCHAR, the maximum length is 32, and the default value is an empty character string.

pv: The type is BIGINT (8 bytes), and the default value is 0. This is an index column, the data will be aggregated by default, and the aggregation method is SUM.

The statement for creating the table is as follows:



```
CREATE TABLE table1
(
  siteid INT DEFAULT '10',
  citycode SMALLINT,
```

```
username VARCHAR(32) DEFAULT '',
pv BIGINT SUM DEFAULT '0'
)
AGGREGATE KEY(siteid,citycode,username)
DISTRIBUTED BY HASH(siteid) BUCKETS 10
PROPERTIES("replication_num" = "1");
```

### Composite partition

Create a logical table named table2, with the following schema:

event\_day: The type is DATE, with no default value.

siteid: The type is INT (4 bytes), and the default value is 10.

citycode: The type is SMALLINT (2 bytes).

username: The type is VARCHAR, the maximum length is 32, and the default value is an empty character string.

pv: The type is BIGINT (8 bytes), and the default value is 0. This is an index column, the data will be aggregated by default, and the aggregation method is SUM.

Use the event\_day column as the partition column and create 3 partitions, including p201706, p201707, and p201708. Each partition uses siteid for hash bucket, and the number of buckets is 10.

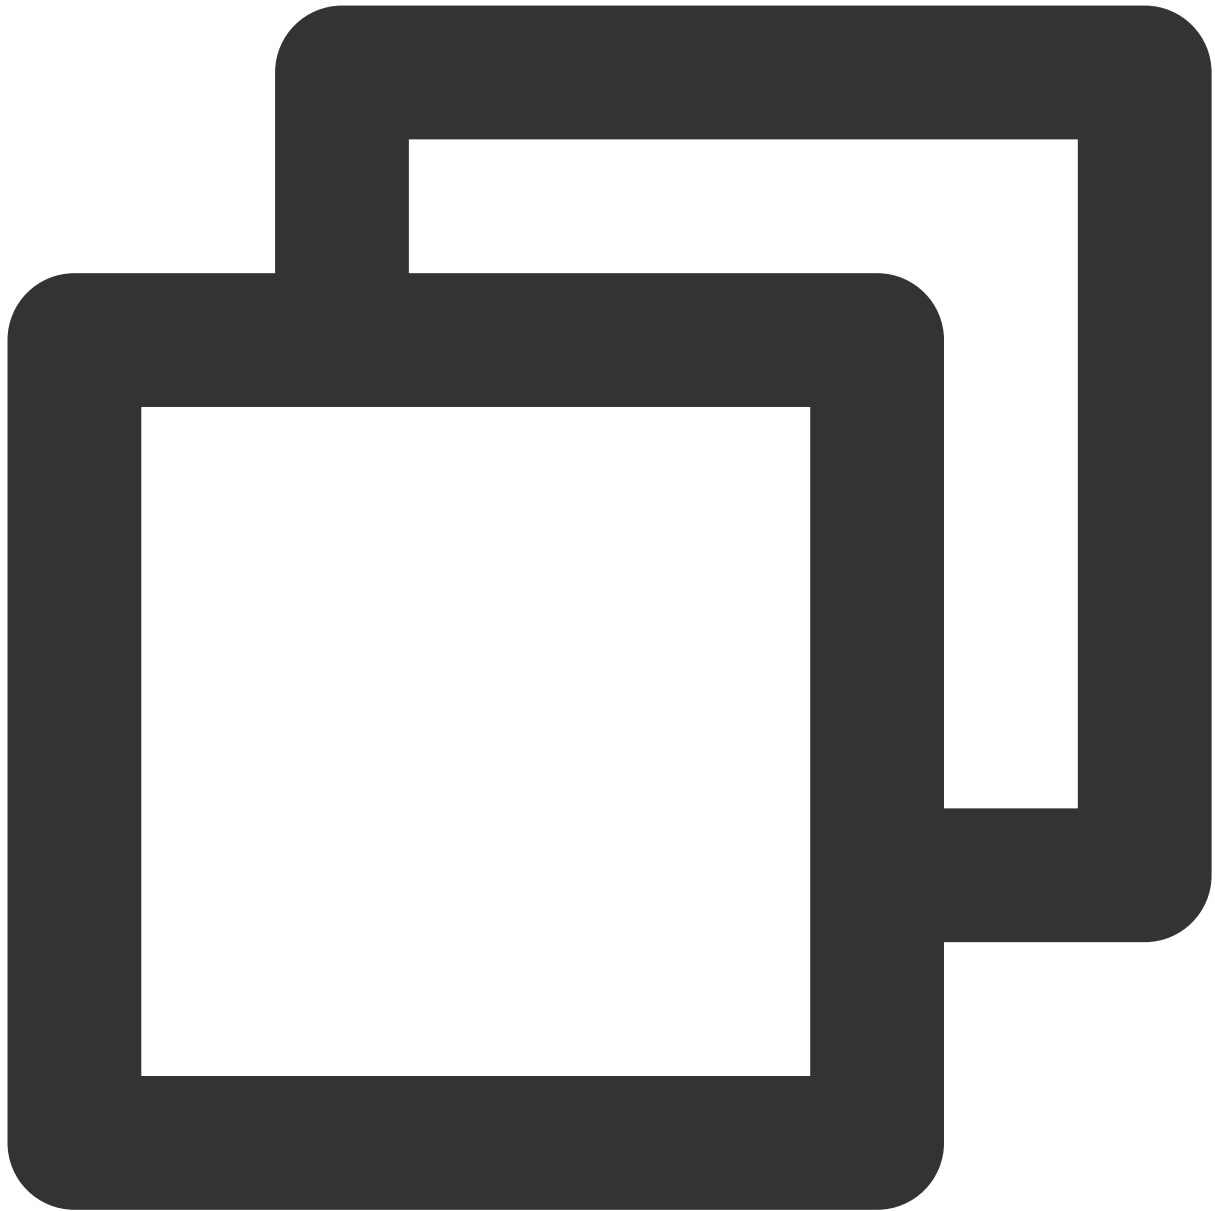
p201706: The range is [minimum value, 2017-07-01).

p201707: The range is [2017-07-01, 2017-08-01).

p201708: The range is [2017-08-01, 2017-09-01).

The statement for creating the table is as follows:






```
CREATE TABLE table2
(
  event_day DATE,
  siteid INT DEFAULT '10',
  citycode SMALLINT,
  username VARCHAR(32) DEFAULT '',
  pv BIGINT SUM DEFAULT '0'
)
AGGREGATE KEY(event_day, siteid, citycode, username)
PARTITION BY RANGE(event_day)
(
```

```
PARTITION p201706 VALUES LESS THAN ('2017-07-01'),  
PARTITION p201707 VALUES LESS THAN ('2017-08-01'),  
PARTITION p201708 VALUES LESS THAN ('2017-09-01')  
)  
DISTRIBUTED BY HASH(siteid) BUCKETS 10  
PROPERTIES("replication_num" = "1");
```

After the table is created, you can check the information of the table in example\_db:



```
MySQL> SHOW TABLES;  
+-----+  
| Tables_in_example_db |
```

```

+-----+
| table1          |
| table2          |
+-----+
2 rows in set (0.01 sec)

MySQL> DESC table1;
+-----+-----+-----+-----+-----+-----+
| Field          | Type          | Null | Key  | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| siteid         | int(11)       | Yes  | true | 10      |       |
| citycode       | smallint(6)  | Yes  | true | N/A     |       |
| username       | varchar(32)  | Yes  | true |         |       |
| pv             | bigint(20)   | Yes  | false | 0       | SUM   |
+-----+-----+-----+-----+-----+-----+
4 rows in set (0.00 sec)

MySQL> DESC table2;
+-----+-----+-----+-----+-----+-----+
| Field          | Type          | Null | Key  | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| event_day      | date          | Yes  | true | N/A     |       |
| siteid         | int(11)       | Yes  | true | 10      |       |
| citycode       | smallint(6)  | Yes  | true | N/A     |       |
| username       | varchar(32)  | Yes  | true |         |       |
| pv             | bigint(20)   | Yes  | false | 0       | SUM   |
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.00 sec)

```

**Note:**

In the above tables, tables created through setting replication\_num are single-copy tables. For Tencent Cloud TCHouse-D, it is recommended that users use the default 3-copy setting to ensure high availability.

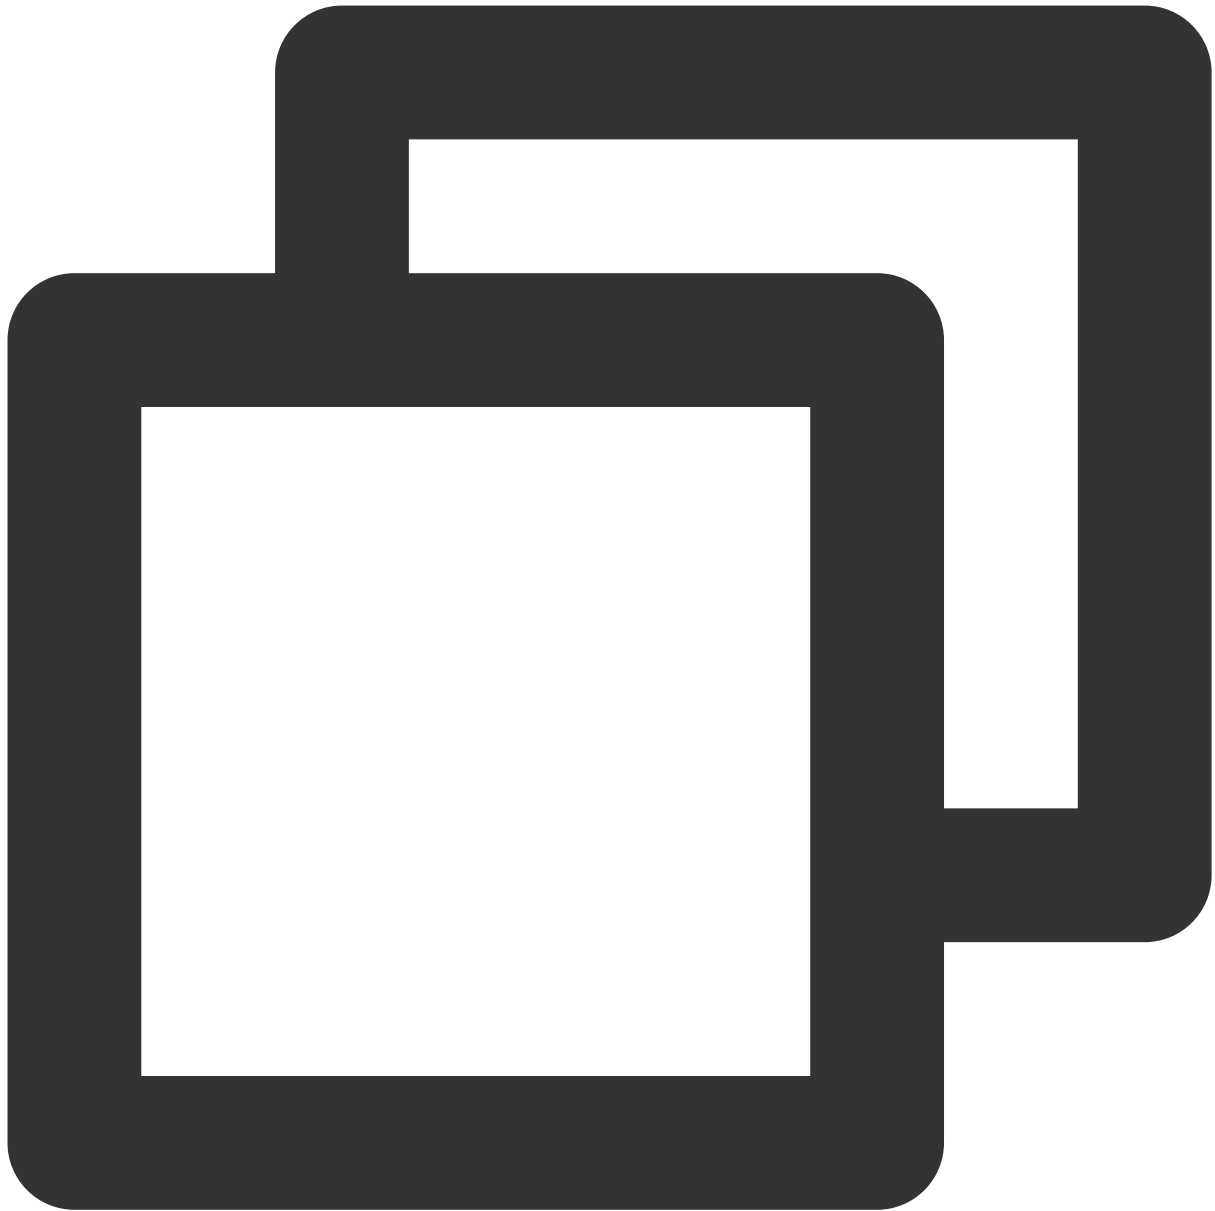
**Step 3: Importing Data**

Tencent Cloud TCHouse-D supports multiple data import methods. For specific details, see [Importing Data](#). Examples of streaming import and Broker import are provided here.

**Streaming Import**

Streaming import transmits data to Tencent Cloud TCHouse-D using the HTTP protocol, allowing for the direct import of local data without dependency on other systems or components. For detailed syntax help, see `HELP STREAM LOAD;`.

Example 1: "table1\_20170707" is used as Label, and the local file table1\_data is imported in table1.



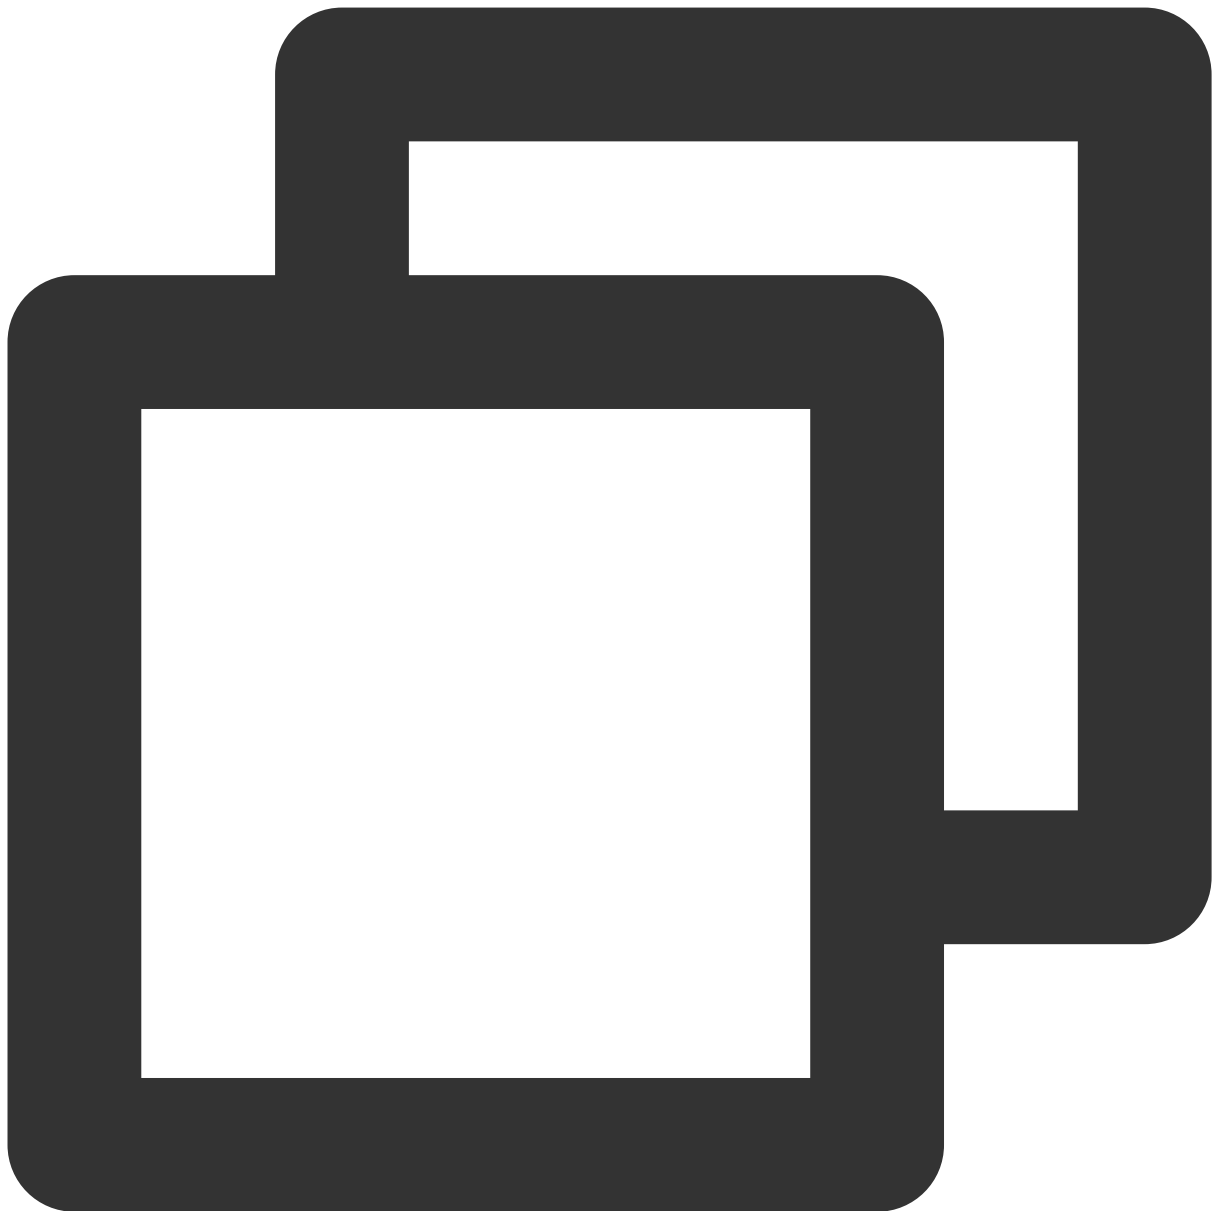
```
curl --location-trusted -u test:test_passwd -H "label:table1_20170707" -H "column_s
```

**Note:**

FE\_HOST is the IP address of any FE's node location, and 8030 is http\_port in fe.conf.

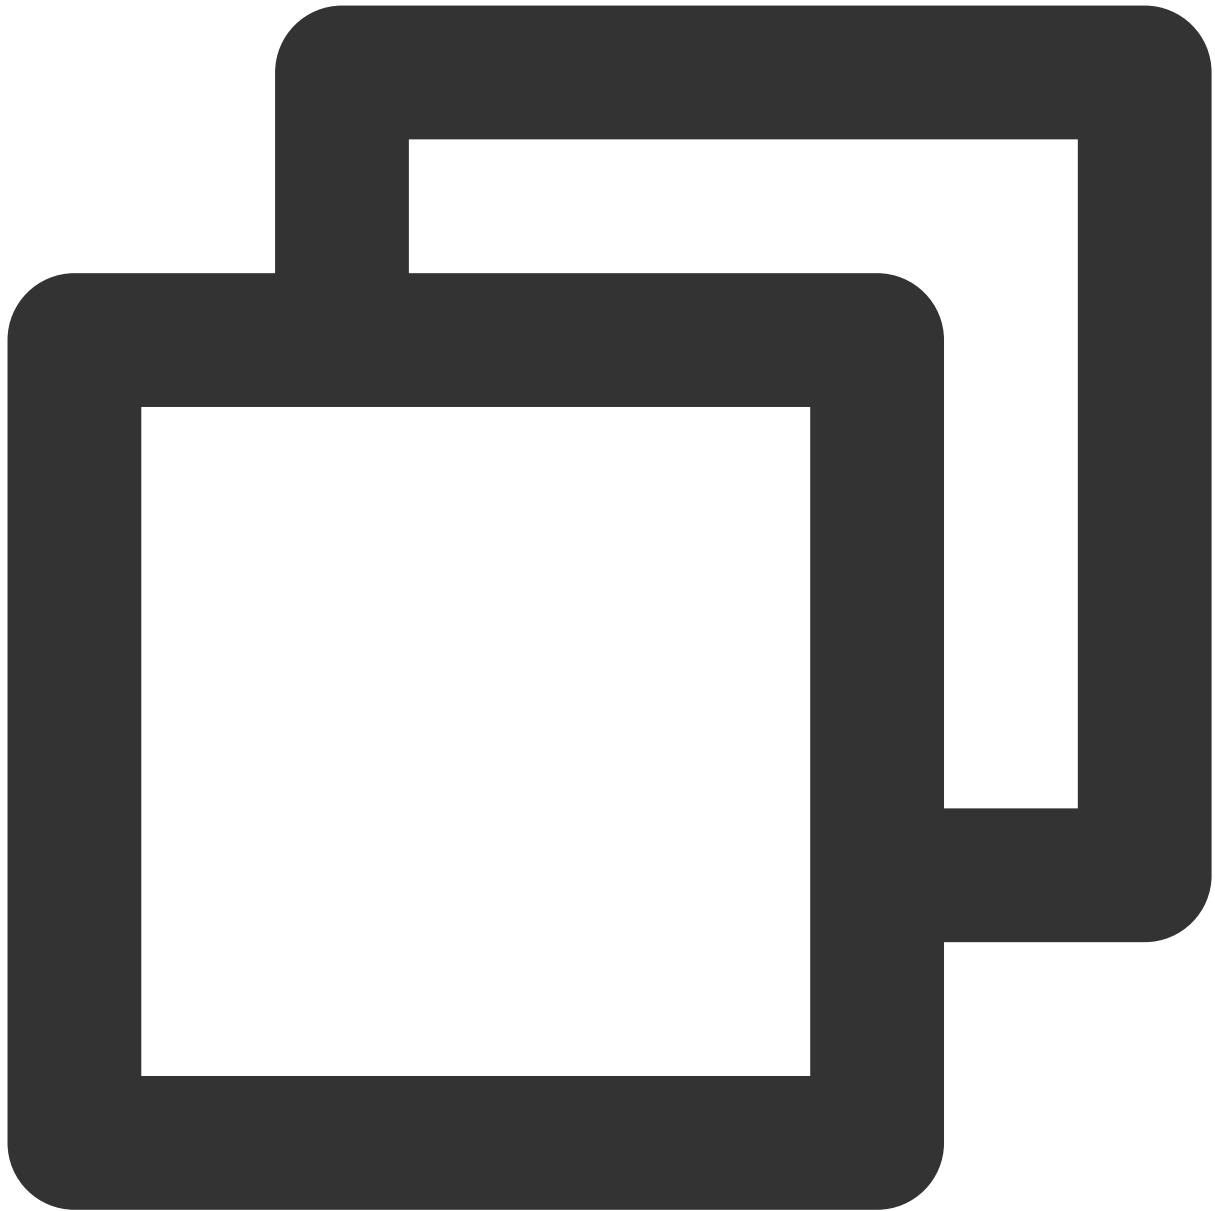
You can use any BE's IP address and webserver\_port in be.conf for importing. Example: BE\_HOST:8040.

For the local file table1\_data, the comma (,) is used as a separator between data, as follows:



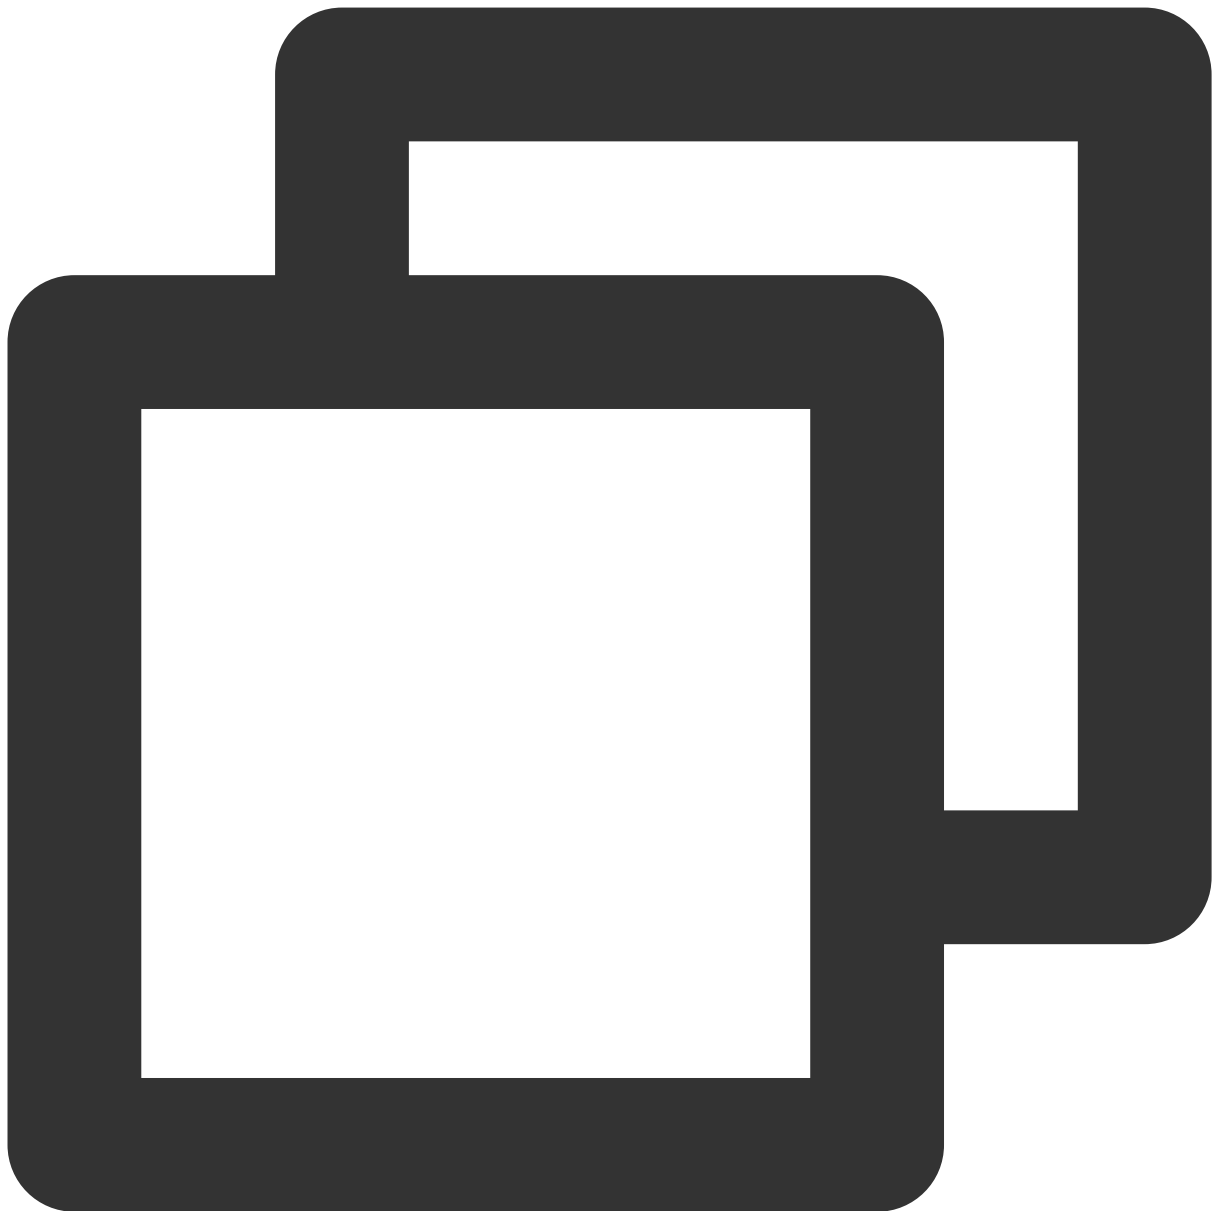
```
1,1,jim,2  
2,1,grace,2  
3,2,tom,2  
4,3,bush,3  
5,3,helen,3
```

Example 2: table2\_20170707 is used as Label, and the local file table2\_data is imported to table2.



```
curl --location-trusted -u test:test -H "label:table2_20170707" -H "column_separato  
http://127.0.0.1:8030/api/example_db/table2/_stream_load
```

For the local file `table2_data`, `|` is used as the separator between data, as follows:



```
2017-07-03|1|1|jim|2  
2017-07-05|2|1|grace|2  
2017-07-12|3|2|tom|2  
2017-07-15|4|3|bush|3  
2017-07-12|5|3|helen|3
```

**Note:**

For streaming import, it is recommended that the file size be within 10 GB. Files that exceed this size could cause import failure and increase the cost of retries.

Each batch of data import requires a unique Label. The Label should ideally be a character string related to the batch of data, facilitating reading and management. Tencent Cloud TCHouse-D guarantees that within a database, the same batch of data can only be imported only once based on the Label. Labels of failed tasks can be reused.

The streaming import command is executed in a synchronous manner. If success is returned for the command, the data has been imported. If failure is returned for the command, the data has not been imported.

### Broker Import

Broker Import uses a deployed Broker process to read and import data from external storage. For more help, see

```
HELP BROKER LOAD; .
```

Example: table1\_20170708 is used as Label, to import files from HDFS into table1.






```
LOAD LABEL table1_20170708
(
  DATA INFILE("hdfs://your.namenode.host:port/dir/table1_data")
  INTO TABLE table1
)
WITH BROKER hdfs
(
  "username"="hdfs_user",
  "password"="hdfs_password"
)
PROPERTIES
```

```
(  
  "timeout"="3600",  
  "max_filter_ratio"="0.1"  
);
```

The Broker import command is executed in an asynchronous manner. The successful execution of the above command only indicates the successful submission of tasks. You can check whether the import is successful through

`SHOW LOAD;` . Example:




```
SHOW LOAD WHERE LABEL = "table1_20170708";
```

In the returned result, if the `State` field is FINISHED, the import is successful. For more information about SHOW LOAD, see `HELP SHOW LOAD;` .

Asynchronous import tasks can be cancelled before they end: `CANCEL LOAD WHERE LABEL = "table1_20170708";` .

## Step 4: Querying Data

### 1. Simple Query:



```
MySQL> SELECT * FROM table1 LIMIT 3;
+-----+-----+-----+
```

```
| siteid | citycode | username | pv |
+-----+-----+-----+-----+
|      2 |         1 | 'grace'  |   2 |
|      5 |         3 | 'helen'  |   3 |
|      3 |         2 | 'tom'    |   2 |
+-----+-----+-----+-----+
3 rows in set (0.01 sec)
```

```
MySQL> SELECT * FROM table1 ORDER BY citycode;
```

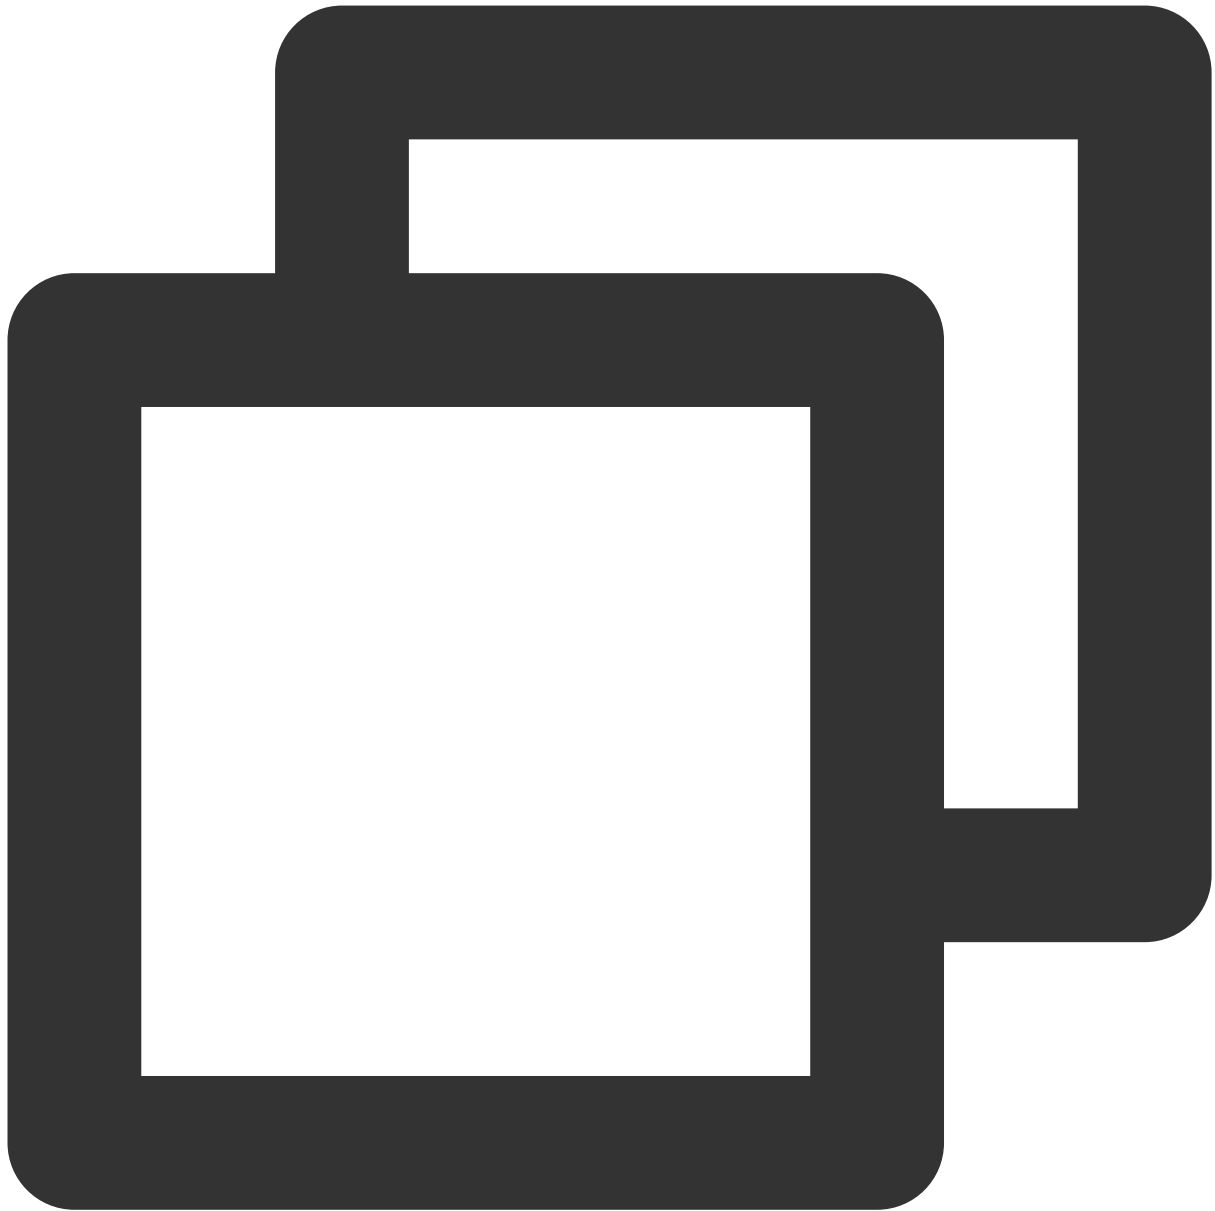
```
+-----+-----+-----+-----+
| siteid | citycode | username | pv |
+-----+-----+-----+-----+
|      2 |         1 | 'grace'  |   2 |
|      1 |         1 | 'jim'    |   2 |
|      3 |         2 | 'tom'    |   2 |
|      4 |         3 | 'bush'   |   3 |
|      5 |         3 | 'helen'  |   3 |
+-----+-----+-----+-----+
5 rows in set (0.01 sec)
```

## 2. Join Query:



```
MySQL> SELECT SUM(table1.pv) FROM table1 JOIN table2 WHERE table1.siteid = table2.s
+-----+
| sum(`table1`.`pv`) |
+-----+
|                12 |
+-----+
1 row in set (0.20 sec)
```

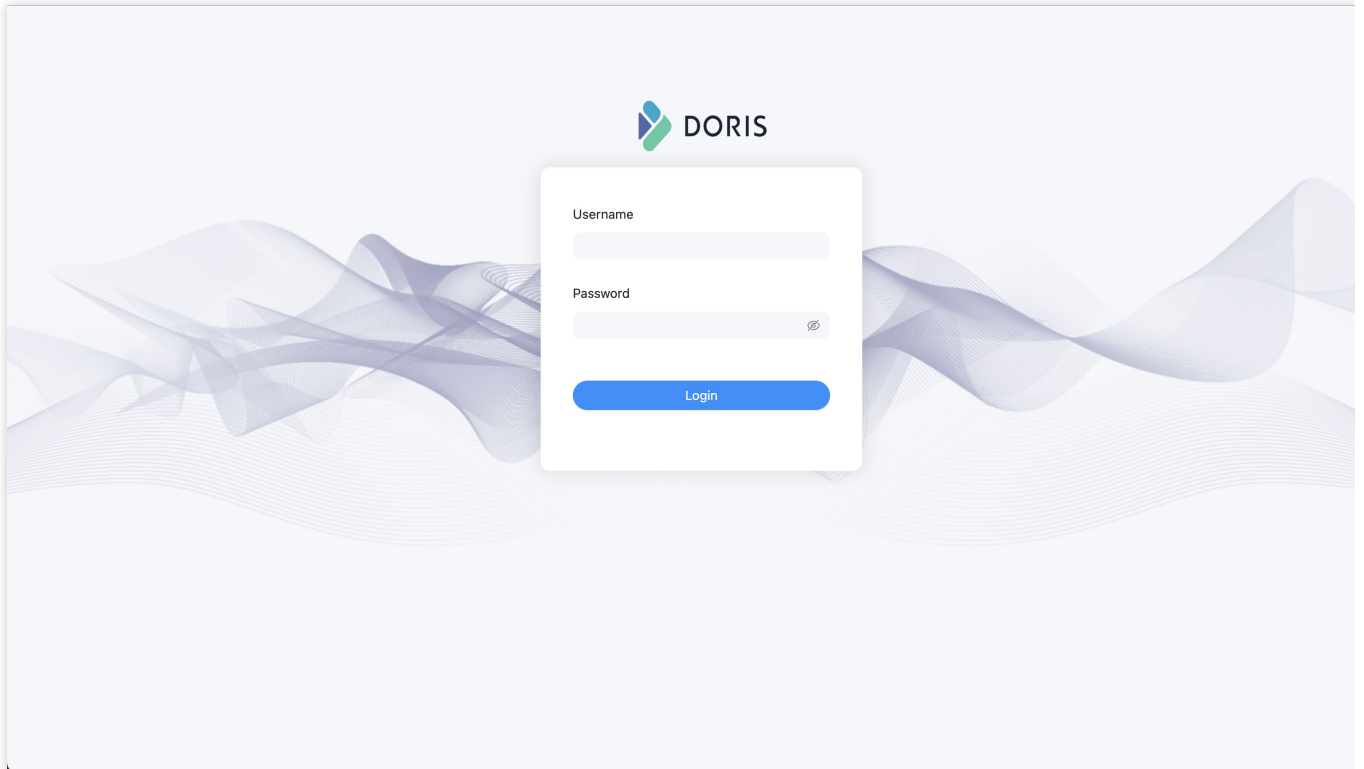
### 3. Subquery:



```
MySQL> SELECT SUM(pv) FROM table2 WHERE siteid IN (SELECT siteid FROM table1 WHERE
+-----+
| sum(`pv`) |
+-----+
|          8 |
+-----+
1 row in set (0.13 sec)
```

## WebUI

The WebUI is based on JDBC and Rest API capabilities, allowing you to perform data operations and manage clusters among other features. To use the WebUI, enter `https:// fe_ip:8030` in a networked environment. The following page is displayed, it has started successfully.



The WebUI is composed of several function pages including Playground, System, Log, QueryProfile, Session, and Configuration.

Playground provides an SQL query editor for executing various SQL commands. The left side displays all databases and their tables in a tree-like structure. You can double-click on a table name to check the metadata and sample data of the table. After you click on `Data Import`, data can be uploaded from local storage.

The screenshot shows the DORIS Playground interface. On the left is a navigation tree with categories like 'default\_cluster:db1', 'default\_cluster:doris\_audit\_db\_', and 'default\_cluster:information\_schema'. The main area contains an 'Editor' with a 'Format' button and a query: `1 select * from db1.t2`. Below the editor is an 'Execute' button. At the bottom, there are tabs for 'Table Schema', 'Data Preview', and 'Data Import'. The 'Data Preview' tab is active, showing a table with columns: TABLE\_CATALOG, TABLE\_SCHEMA, TABLE\_NAME, COLUMN\_NAME, ORDINAL\_POSITION, and COLUMN\_DEFA. The table contains four rows of data for the 'partitions' table in the 'information\_schema' database.

TABLE_CATALOG	TABLE_SCHEMA	TABLE_NAME	COLUMN_NAME	ORDINAL_POSITION	COLUMN_DEFA
	information_schema	partitions	TABLE_CATALOG	1	
	information_schema	partitions	TABLE_SCHEMA	2	
	information_schema	partitions	TABLE_NAME	3	
	information_schema	partitions	PARTITION_NAME	4	

The System page provides various system information about the cluster.



**DORIS** Playground **System** Log QueryProfile Session Configuration

## System Info

This page lists the system info, like `/proc` in Linux.

Current path: `/`

name
<a href="#">auth</a>
<a href="#">backends</a>
<a href="#">bdbje</a>
<a href="#">brokers</a>
<a href="#">cluster_balance</a>
<a href="#">cluster_health</a>
<a href="#">colocation_group</a>
<a href="#">current_backend_instances</a>
<a href="#">current_queries</a>
<a href="#">current_query_stmts</a>

The Log page provides features for displaying and managing the FE logs.

The screenshot shows the DORIS web interface. At the top, there are navigation tabs: Playground, System, Log (selected), QueryProfile, Session, and Configuration. The main heading is "Log Configuration". Below it, there are fields for "Level:", "Verbose Names:", and "Audit Names: slow\_query,query,load,stream\_load". There are two input boxes: "new verbose name" with an "Add" button, and "del verbose name" with a "Delete" button. Below this is the "Log Contents" section, showing the "Log path is: /data/cdw/doris/fe/log/fe.warn.log" and "38334 bytes of log". A code block displays a stack trace for a failed stmt execution.

```

2022-09-28 21:46:58,704 WARN (qtp1286393023-16618|16618) [StmtExecutionAction.executeQuery():154] failed to execute stmt
java.util.concurrent.ExecutionException: java.sql.SQLException: Socket error (conn=81)
    at java.util.concurrent.FutureTask.report(FutureTask.java:122) ~[?:?]
    at java.util.concurrent.FutureTask.get(FutureTask.java:191) ~[?:?]
    at org.apache.doris.httpv2.rest.StmtExecutionAction.executeQuery(StmtExecutionAction.java:148) ~[doris-fe.jar:1.0-SNAPSHOT]
    at org.apache.doris.httpv2.rest.StmtExecutionAction.executeSQL(StmtExecutionAction.java:103) ~[doris-fe.jar:1.0-SNAPSHOT]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke0(Native Method) ~[?:?]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) ~[?:?]
    at jdk.internal.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) ~[?:?]
    at java.lang.reflect.Method.invoke(Method.java:566) ~[?:?]
    at org.springframework.web.method.support.InvocableHandlerMethod.doInvoke(InvocableHandlerMethod.java:205) ~[spring-web-5.3.22.jar:5.3.22]
    at org.springframework.web.method.support.InvocableHandlerMethod.invokeForRequest(InvocableHandlerMethod.java:150) ~[spring-web-5.3.22.jar:5.3.22]
    
```

The QueryProfile page displays SQL recorded in the profile, where you can check the overall SQL as well as the execution plan and runtime data for each instance.

The screenshot shows the DORIS web interface with the "QueryProfile" tab selected. The heading is "Finished Queries". Below it, a note says "This table lists the latest 100 queries". A table with columns: Detail, Job ID, Query ID, User, Default Db, and Sql Statement. Two rows are visible:

Detail	Job ID	Query ID	User	Default Db	Sql Statement
fc25e5792ac54798-8c08d39ef9a1bf0d	N/A	fc25e5792ac54798-8c08d39ef9a1bf0d	root	default_cluster:tpch_100g_11	SELECT DATABASE0
2754d3a095534783-bab9937e471fc306	N/A	2754d3a095534783-bab9937e471fc306	root	default_cluster:tpch_100g_11	select l_returnflag, l_linestatus, sum(l_quantity) as sum_qty, sum(l_extendedprice) as sum_base_price, sum(l_extendedpr

The Session page displays information about all active sessions.

**DORIS** Playground System Log QueryProfile Session Configuration

## Session Info

This page lists the session info, there are 1 active sessions.

Id	User	Host	Cluster	Db	Command
80	admin	127.0.0.1:56628	default_cluster		Sleep

The Configuration page provides all configuration information of the FE, offering a filtering feature in the configuration value column.

**DORIS** Playground System Log QueryProfile Session Configuration

## Configure Info

Name
sys_log_roll_interval
max_routine_load_job_num
query_colocate_join_memory_limit_penalty_factor
remote_fragment_exec_timeout_ms
max_running_rollup_job_num_per_table
max_allowed_in_element_num_of_delete
async_loading_load_task_pool_size
max_stream_load_timeout_second
max_routine_load_task_num_per_be
using_old_load_usage_pattern

1-10 of 253 items