

# CloudTable Service

## FAQs

**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.

# 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

---

<b>1 General.....</b>	<b>1</b>
1.1 What Services Does a CloudTable Cluster Provide?.....	1
1.2 Why Do I Choose CloudTable Service?.....	1
1.3 How Do I Prepare for Creating a CloudTable HBase Cluster?.....	1
1.4 What Should I Pay Attention to When Using CloudTable Service?.....	2
1.5 What Compression Algorithms Are Supported by CloudTable HBase Clusters?.....	2
1.6 Can I Stop CloudTable Service?.....	2
1.7 Which Programming Languages Are Supported by HBase External APIs in CloudTable?.....	2
1.8 How Do I Determine the Number of Faulty RegionServers?.....	3
1.9 What Special Characters Does CloudTable HBase Support?.....	3
1.10 What Can I Do If the Index Table Does Not Match the Data Table After CloudTable Data Is Deleted?.....	3
1.11 What Should I Do If Concurrent Tasks Run Slowly When Python Accesses CloudTable Through Thrift?.....	3
1.12 How do I view the TTL attribute of HBase shell?.....	4
1.13 Why Are My Server Resources Released?.....	4
1.14 How Do I Delete a Cluster?.....	4
1.15 How Do I Stop Services and Release Resources?.....	5
<b>2 Connection and Access.....</b>	<b>6</b>
2.1 How Do I Access a CloudTable Cluster?.....	6
2.2 Can I Use SSH to Access Computing Nodes of CloudTable?.....	6
2.3 Why Can't I Access HBase After the ZK Address Is Configured?.....	6
2.4 Why Is the Error "Will not attempt to authenticate using SASL (unknown error)" Reported When Connecting to HBase?.....	7
2.5 How Do I View the IP Address Corresponding to a Domain Name in a CloudTable Link?.....	7
2.6 How Do I Access CloudTable from Other Cloud Services?.....	7
2.7 Can I Configure the hbase-site.xml File?.....	8
2.8 How Do I Query the Creation Time of a Table in CloudTable HBase?.....	8
<b>3 Data Read/Write.....</b>	<b>9</b>
3.1 Is Raw Data Stored in CloudTable HBase?.....	9
3.2 Why Can't I Write Data to HBase?.....	9
3.3 What Is the Maximum Size of Data Written to the HBase Cluster?.....	10
3.4 How Do I Check the Daily Incremental Data in HBase Tables?.....	10

---

3.5 What Should I Do If an Error Is Reported When I Access the CloudTable HBase Cluster?.....	10
3.6 How Do I Delete the Backup Table of the ZooKeeper Node in the ClickHouse Cluster?.....	11
3.7 What Should I Do If a Database Missing Error Occurs When a Table Is Created in the ClickHouse Cluster?.....	12
<b>4 Billing FAQs.....</b>	<b>14</b>
4.1 How Do I Renew a CloudTable Cluster?.....	14
4.2 What Is Billed Duration?.....	15

# 1 General

---

## 1.1 What Services Does a CloudTable Cluster Provide?

CloudTable Service (CloudTable) provides you with dedicated clusters, which are available upon subscription. It is suitable for users with high service throughput and low delay requirements.

- CloudTable offers HBase-based fully managed NoSQL services, delivering lightning-fast random read/write capabilities within milliseconds. It is an excellent choice for storing vast amounts of structured, semi-structured, spatiotemporal, and time series data. Whether in IoT, IoV, finance, smart cities, or meteorology, CloudTable finds versatile applications.
- ClickHouse is an open-source columnar database oriented to online analysis and processing. It is independent of the Hadoop big data system and features compression rate and fast query performance. In addition, ClickHouse supports SQL query and provides good query performance, especially the aggregation analysis and query performance based on large and wide tables. The query speed is one order of magnitude faster than that of other analytical databases.

## 1.2 Why Do I Choose CloudTable Service?

CloudTable is a fully managed service based on the public cloud. It provides you with dedicated clusters, which are available upon subscription. It is suitable for users who require high service throughput and low latency. CloudTable is O&M-free and features effortless deployment. It allows you to create clusters on demand to reduce costs and offers professional tutorial.

## 1.3 How Do I Prepare for Creating a CloudTable HBase Cluster?

- Select a proper region to create a CloudTable cluster based on cluster access delay requirements.

- When creating a CloudTable cluster, determine the number of computing units based on service loads. After a cluster is created, computing units cannot be scaled-in and the cluster is charged based on the number of computing units. If the delay for the cluster to process services exceeds the recommended value, increase the number of computing units.

## 1.4 What Should I Pay Attention to When Using CloudTable Service?

- On the CloudTable management console, check cluster status and task status. If the cluster status is **Creation failed** or **Sub-health**, or a task fails to be executed, submit a service ticket to technical support.
- On the CloudTable management console, click **Monitor** on the right of the **Cluster Management** page to access the **CloudTable Service Monitoring** page on the Cloud Eye management console. View metric monitoring information and set alarm rules for key metrics. If there is an alarm, submit a service ticket to technical support. If the cluster CPU or memory usage exceeds the threshold for a long time, you need to increase the number of computing units to balance loads.
- On the CloudTable console, click **Cluster Management** to view the cluster list. Check whether the used storage capacity of a cluster exceeds the threshold. If it exceeds the threshold, increase storage quotas and storage capacities.

## 1.5 What Compression Algorithms Are Supported by CloudTable HBase Clusters?

CloudTable HBase clusters support the SNAPPY and GZ compression algorithms.

## 1.6 Can I Stop CloudTable Service?

No, you cannot. CloudTable Service complies with the common practices in the industry. Once a cluster is created, it cannot be stopped or shut down.

## 1.7 Which Programming Languages Are Supported by HBase External APIs in CloudTable?

Currently, external APIs of HBase 2.4.14 only support the Java programming language.

For more information about HBase APIs, see [Apache HBase External APIs](#) in the HBase official documentation.

## 1.8 How Do I Determine the Number of Faulty RegionServers?

Number of faulty RegionServer instances in the cluster where the monitored object is located.

## 1.9 What Special Characters Does CloudTable HBase Support?

There is no restriction on special characters in service data supported by HBase. Only colons (:) in Lemon need to be escaped.

## 1.10 What Can I Do If the Index Table Does Not Match the Data Table After CloudTable Data Is Deleted?

### Symptom

When you query data in the CloudTable table, the value of **count** is inconsistent with the number of data records in **list**. Data loss frequently occurs, and many links time out during project running.

### Possible Causes

After data is deleted, the index table does not match the data table. As a result, an exception occurs.

### Solutions

- Short-term solution: Perform offline batch processing analysis on data streams and supplement the missing data on the next day.
- Long-term solutions:
  - The lemon feature does not support data deletion. After data is deleted, table data cannot be queried. In this case, disable the TTL aging function of the data table.
  - Check whether the TTL enhanced index can be removed in the CBC data aging scenario.

## 1.11 What Should I Do If Concurrent Tasks Run Slowly When Python Accesses CloudTable Through Thrift?

### Symptom

When a CloudTable HBase data table is accessed, concurrent tasks run slowly.



## Possible Causes

GC memory leakage occurs on cluster nodes. Perform gray restart on the nodes.

### NOTE

In addition, check whether the load is too high. If the nodes are overloaded, contact related personnel.

## Solutions

Contact maintenance personnel to restart the RegionServer instance.

## 1.12 How do I view the TTL attribute of HBase shell?

- If the TTL of the cell is set when the data is inserted, the TTL attribute cannot be viewed. You can check whether the TTL configuration is successful.
- If the TTL of the cell is not set when data is inserted, the system automatically inserts the current time as the timestamp.

## 1.13 Why Are My Server Resources Released?

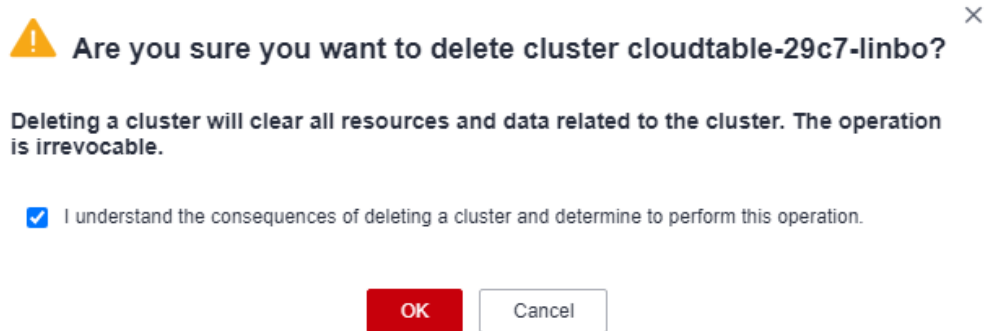
After you purchase a resource on Huawei Cloud but you do not complete the payment or renewal for the product, you will enter a grace period. If you still do not complete the payment or renewal after the grace period has ended, you will enter a retention period, during which the resources will be suspended. If you still do not complete the payment or renewal after the retention period has ended, your data stored in the cloud service will be deleted and the resource will be released.

## 1.14 How Do I Delete a Cluster?

You can delete clusters you will no longer use.

- Step 1** Log in to the console and go to the **Cluster Management** page.
- Step 2** Locate the row that contains the cluster to be deleted and click **Delete** in the **Operation** column.
- Step 3** Read the displayed message and confirm the deletion.

**Figure 1-1** Confirming the deletion



**Step 4** Select "I understand the consequences of deleting a cluster and determine to perform this operation." and click **OK**.

----End

## 1.15 How Do I Stop Services and Release Resources?

After you purchase a resource on Huawei Cloud but you do not complete the payment or renewal for the product, you will enter a grace period. If you still do not complete the payment or renewal after the grace period has ended, you will enter a retention period, during which the resources will be suspended. If you still do not complete the payment or renewal after the retention period has ended, your data stored in the cloud service will be deleted and the resource will be released.

- **Grace period:** When you do not renew a yearly/monthly subscription before it expires or a pay-per-use resource is in the arrears state, Huawei Cloud provides a grace period during which you can renew the resource or top up your account. During the grace period, you can still access and use your cloud service. Huawei Cloud will send you an email or text message to remind you to renew the resource or top up your account.
- **Retention period:** During the grace period, if you do not renew the yearly/monthly subscription or pay off the arrears for the pay-per-use resource, the resource enters the retention period after the grace period has expired. During this period, your cloud services cannot be accessed or used, but your stored cloud data will be retained. Huawei Cloud will send you an email or text message to remind you to renew the resource or top up your account.

# 2 Connection and Access

---

## 2.1 How Do I Access a CloudTable Cluster?

CloudTable is based on an elastic cloud server (ECS). You need to create an ECS in the same subnet as the CloudTable cluster, synchronize the ZooKeeper address of the cluster, and access the cluster using the client or APIs.

## 2.2 Can I Use SSH to Access Computing Nodes of CloudTable?

No. Underlying computing units of CloudTable run on a Virtual Private Cloud (VPC) network. VPC isolates networks logically and provides a secure and isolated network environment to protect your data. You can use only a client tool or application to connect to the ZooKeeper link (that is, an HBase connection address) or OpenTSDB link of CloudTable to access CloudTable.

 NOTE

The OpenTSDB function applies only to HBase 1.3.1.

## 2.3 Why Can't I Access HBase After the ZK Address Is Configured?

CloudTable does not provide an external access address. Therefore, you cannot access HBase from the public network. You need to prepare an ECS that is in the same VPC as the HBase cluster to access HBase.

Ensure that **nameserver** is correctly configured in the `/etc/resolv.conf` file on your ECS. For details about configuration methods, see [Preparing an ECS](#) in the *CloudTable Service User Guide*.

## 2.4 Why Is the Error "Will not attempt to authenticate using SASL (unknown error)" Reported When Connecting to HBase?

### Symptom

HBase fails to be connected to when accessing CloudTable and the following error information is displayed:

```
Opening socket connection to server 192.168.0.107/192.168.0.107:2181. Will not attempt to authenticate using SASL (unknown error)
```

### Possible Cause

The possible cause is as follows:

Network access failure

The link of CloudTable is an internal network address rather than a public network address. Therefore, CloudTable cannot be directly accessed from the public network. You can ping the ZooKeeper link of CloudTable on a client host. If it cannot be pinged, the network connection is abnormal.

### Solution

You need to apply for an ECS as the client host to install the CloudTable client. Then, you can use the client on the ECS to access the CloudTable cluster. The ECS to be applied for and the CloudTable cluster must be in the same region, VPC, subnet, and security group. In addition, you need to configure a correct DNS server address for the ECS. For details about how to prepare an ECS, see [Preparing an ECS](#).

## 2.5 How Do I View the IP Address Corresponding to a Domain Name in a CloudTable Link?

Submit a service ticket to technical support to obtain the IP address corresponding to a domain name, because you cannot view it by yourself.

To submit a service ticket, log in to the CloudTable management console, and choose **More > Service Tickets > Create Service Ticket** in the upper right corner. Fill in the service ticket to obtain the mapping between the intranet domain name of the CloudTable link and the IP address.

## 2.6 How Do I Access CloudTable from Other Cloud Services?

When a cloud service and a CloudTable cluster are in the same region, AZ, VPC, subnet, and security group, you can access the CloudTable from the cloud service.

## 2.7 Can I Configure the hbase-site.xml File?

- hbase-site.xml file on a client  
If you have deployed an HBase shell client, you can configure the **hbase/conf/hbase-site.xml** file in the client directory.
- hbase-site.xml file on a server  
You cannot edit the **hbase-site.xml** file on the server. However, you can modify some parameters in the **hbase-site.xml** file on the CloudTable management console. For details, see [Modifying HBase Parameters of the Cluster](#). If modification of these parameters cannot meet your requirements and you want to modify other parameters, contact technical support.

## 2.8 How Do I Query the Creation Time of a Table in CloudTable HBase?

You can use either of the following methods to query the table creation time in CloudTable:

Method 1: Check the table creation time in the output of the **create** command.

**Step 1** Connect to the cluster and use the HBase shell to access the cluster. For details, see [Using HBase Shell to Access a Cluster](#) in the *CloudTable Service User Guide*.

**Step 2** Run the **create** command to create a table.

```
create 'TEST','StuInfo','Grades'
```

**Step 3** In the command output, the first row indicates the table creation time and table name, and the second row indicates the table creation duration.

**Figure 2-1** Table creation time

```
hbase(main):004:0> create 'TEST','StuInfo','Grades'  
2023-03-22 11:17:34,351 INFO [main] client.HBaseAdmin: Created TEST  
0 row(s) in 1.2210 seconds
```

----End

# 3 Data Read/Write

---

## 3.1 Is Raw Data Stored in CloudTable HBase?

CloudTable HBase stores raw data and supports compressed storage. The SNAPPY and GZ compression algorithms are supported. CloudTable HBase does not encrypt or decrypt raw data. If data needs to be encrypted or decrypted, use service code on an application to implement data encryption and decryption.

## 3.2 Why Can't I Write Data to HBase?

### Symptom

Data cannot be written to HBase. As a result, service data on an application cannot be updated in time.

### Possible Cause

There are too many ZooKeeper connections.

### Solution

When service code is connected to the same CloudTable HBase cluster, you are advised to create one Connection and reuse it for multiple threads. You do not need to create a Connection for every thread. Connection is a connector for connecting to a CloudTable HBase cluster. Excessive Connections will increase loads on ZooKeeper and deteriorate service read/write performance. Multiple client threads can share one Connection. Typically, client threads of a client program share one Connection.

## 3.3 What Is the Maximum Size of Data Written to the HBase Cluster?

### Symptom

When a large number of data is concurrently written to an HBase cluster, data cannot be fully written.

### Possible Causes

The data written at a time is too large and an error occurs.

### Solutions

The size of data written at a time should not exceed 2 MB and the size of a single data record should not exceed 200 KB.

## 3.4 How Do I Check the Daily Incremental Data in HBase Tables?

Incremental data in HBase tables cannot be directly checked. You need to query the incremental data in the background. Alternatively, you can view the daily incremental data on the console, view the used storage capacity on the details page, and calculate the incremental data by recording the storage capacity.

## 3.5 What Should I Do If an Error Is Reported When I Access the CloudTable HBase Cluster?

### Symptom

An error is reported when the CloudTable HBase cluster is accessed, indicating that the hmaster domain name cannot be identified.

### Possible Causes

The mapping between the private domain name and IP address of the CloudTable HBase instance nodes is not configured on the server where the program is running. As a result, the private domain name of the hmaster node cannot be identified, and a link error occurs.

### Solutions

In the **host** file on the server where the program is running, configure the mapping between the private domain name and IP address of CloudTable HBase instance nodes. For details, see "[Preparing an ECS](#)".

## 3.6 How Do I Delete the Backup Table of the ZooKeeper Node in the ClickHouse Cluster?

### Symptom

When you connect a ClickHouse cluster and create a table, the ZooKeeper node does not get fully removed after you delete the table.

**Figure 3-1** Execution results

```
ENGINE = MergeTree
ORDER BY id
Query id: 9765507e-bd39-4409-9810-1d8ca82fda4d
```

host	port	status	error	num_hosts_remaining	num_hosts_active
192.168.5.11	9000	0		1	0
192.168.5.117	9000	57	Code: 57. DB::Exception: Table tt.table1 already exists. (TABLE_ALREADY_EXISTS) (version 22.3.2.1)	0	0

Progress: 0.00 rows, 0.00 B (0.00 rows/s., 0.00 B/s.) 0%  
2 rows in set. Elapsed: 0.113 sec.

### Possible Causes

The command for creating a table contains **ON CLUSTER default\_cluster**. Therefore, a local table is created on each node.

### Solutions

Delete the tables and stored data on other nodes.

**Step 1** Create a database.

```
create database demo ON CLUSTER default_cluster;
```

**Step 2** Create tables in the database.

```
use demo;
```

Create the **test** table.

```
CREATE TABLE demo.test ON CLUSTER default_cluster('EventDate` DateTime, `id` UInt64)ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/test', '{replica}') PARTITION BY toYYYYMM(EventDate) ORDER BY id;
```

**Step 3** Delete the table.

```
drop table test SYNC;
```

**Step 4** The following error message is displayed if a table with the same name as the deleted one is created:

```
host | port | status | error | num_hosts_remaining | num_hosts_active
```

host	port	status	error	num_hosts_remaining	num_hosts_active
192.168.2.185	9000	57	Code: 57. DB::Exception: Table demo.test already exists. (TABLE_ALREADY_EXISTS) (version 22.3.2.1)	1	1

host	port	status	error	num_hosts_remaining	num_hosts_active
192.168.2.16	9000	0		0	0

**Step 5** Method 1: Run the following command to delete the tables and related data on each node:



```
drop table test ON CLUSTER default_cluster SYNC;
```

No error occurs if a table with the same name as the deleted one is created.

**Step 6** Method 2: Delete the table on other nodes.

```
drop table test SYNC;
```

----End

## 3.7 What Should I Do If a Database Missing Error Occurs When a Table Is Created in the ClickHouse Cluster?

### Symptom

When a table is created, the system displays a message indicating that the database does not exist on another node.

**Figure 3-2** Execution results

host	port	status	error	num_hosts_remaining	num_hosts_active
192.168.2.16	9000	81	Code: 81. DB::Exception: Database 'demo' doesn't exist. (UNKNOWN_DATABASE) (version 22.3.2.1)	1	0
192.168.2.185	9000	0		0	0

### Possible Causes

A database is created only on one node when the cluster is created.

### Solutions

Create a database on another code or use another command to create one on each node.

Method 1:

**Step 1** Creates a database.

```
create database demo;
```

Use the **demo** database.

```
use demo;
```

**Step 2** Create the **test** table. The result shown in [Figure 1](#) is displayed. The **demo** database does not exist on other nodes.

```
CREATE TABLE demo.test ON CLUSTER default_cluster('EventDate` DateTime, `id` UInt64)ENGINE = ReplicatedMergeTree('/clickhouse/tables/{shard}/default/test', '{replica}') PARTITION BY toYYYYMM(EventDate) ORDER BY id;
```

**Step 3** Create a database **demo** on another node and create a table. If the information shown in the following figure is displayed, no error occurs.

**Figure 3-3** Successful execution

host	port	status	error	num_hosts_remaining	num_hosts_active
192.168.2.185	9000	0		1	1
192.168.2.16	9000	0		0	0

----End

Method 2:

**Step 1** Delete the existing database and use the following command to create a database on each node.

```
create database demo ON CLUSTER default_cluster;
```

**Step 2** Create a table and error message shown in [Figure 1](#) is not displayed.

----End

# 4 Billing FAQs

---

## 4.1 How Do I Renew a CloudTable Cluster?

### Renewal Description

Currently, CloudTable clusters support the pay-per-use and yearly/monthly billing modes.

- In pay-per-use (hourly) mode, the system settles the fees by hour. You can use CloudTable as long as your account balance is sufficient.
- In yearly/monthly mode, you pay only once when purchasing the service and no extra fees will be charged during your use of the service. After a yearly/monthly subscription expires, the resources will enter a retention period. If you want to continue using the resources, renew the subscription. For details, check the resource expiration rules.
- If your account balance is insufficient and fee deduction fails, a retention period starts. During the retention period, the service resources are frozen and cannot be used, but resources and data are reserved.
- If your account is not recharged within the retention period, the frozen service resources and data will be deleted and cannot be restored after the retention period expires.
- You can recharge your account within the retention period. After the recharge is successful, the system preferentially deducts the arrears and unfreezes the service resources. Then you can continue to use CloudTable properly.

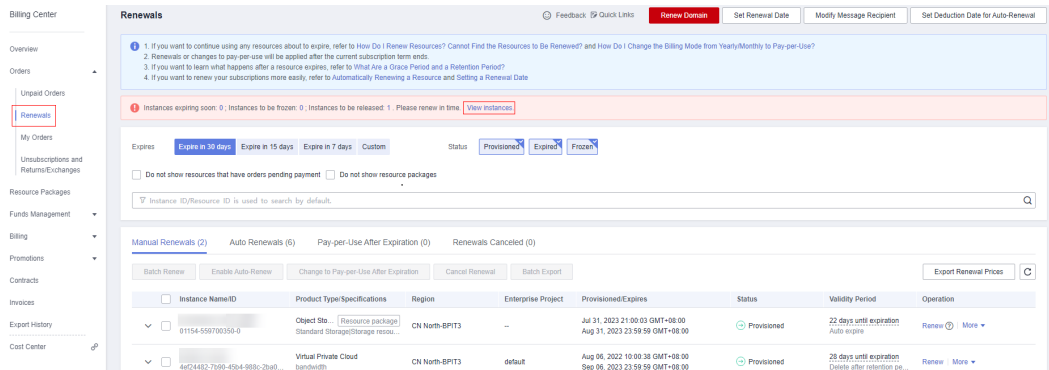
Perform the following operations to recharge your account by yourself to ensure that you can use CloudTable properly.

### How to Renew the Service

**Step 1** Log in to the CloudTable console.

**Step 2** Choose **Billing > Renewal** in the upper right corner of the page. The **Renewal** page is displayed.

Figure 4-1 Viewing resources



**Step 3** On the **Renewals** page, choose the renewal mode based on your service requirements.

**Step 4** Complete the renewal as prompted.

- You cannot change the billing mode of an order during renewal.
- After you pay for a new order, it takes effect only after the old order expires.

----End

## 4.2 What Is Billed Duration?

The billed duration begins when you purchase a cluster and ends when you delete it.