

CloudTable Service

User Guide

Issue 19
Date 2019-11-07

Copyright © Huawei Technologies Co., Ltd. 2020. 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

Contents

1 Getting Started (Cluster Mode)	1
1.1 Creating a Cluster Quickly	1
1.2 Getting Started with HBase.....	3
2 Permissions Management	6
2.1 Creating a User and Granting Permissions	6
3 Cluster Mode	9
3.1 Introduction to the Cluster Mode.....	9
3.2 Managing Clusters	10
3.2.1 Creating a Cluster	10
3.2.2 Purchasing a Discount Package	13
3.2.3 Introduction to Cluster Management.....	15
3.2.4 Viewing Basic Cluster Information	17
3.2.5 Modifying HBase Parameters of the Cluster	19
3.2.6 Enabling OpenTSDB.....	23
3.2.7 Restarting a Cluster	23
3.2.8 Expanding Cluster Capacity.....	24
3.2.9 Deleting a Cluster	26
3.3 Preparing a Client Operating Environment.....	26
3.3.1 Preparing an ECS.....	26
3.3.2 Preparing a Local Windows Environment (VPN Connection Mode)	30
3.4 Using HBase.....	39
3.4.1 Using HBase Shell to Access a Cluster	39
3.4.2 HBase Shell Commands.....	44
3.4.3 Accessing the HBase Web UI	45
3.4.4 Authorizing Other Users to Access HBase in a Cluster with IAM Authentication Enabled	49
3.5 Using OpenTSDB	51
3.5.1 OpenTSDB Overview	51
3.5.2 Connecting to OpenTSDB	53
3.6 Using GeoMesa.....	54
3.6.1 Using the GeoMesa Command Line Tools	54
3.6.2 GeoMesa Command Line	56
3.7 HBase Elasticsearch Full-Text Search.....	58

3.7.1 Overview of Full-Text Search.....	58
3.7.2 Using the HBase Shell for Full-Text Indexing	61
3.8 Batch Data Import.....	62
3.8.1 Using CDM to Migrate Data to CloudTable	62
3.8.2 Using the Import Tool to Importing Data.....	66
3.8.3 Using CopyTable to Import Data.....	66
3.8.4 Sample Code for Copying a Table	68
3.9 Cluster Monitoring.....	71
3.9.1 Monitored CloudTable Metrics	71
3.9.2 Setting CloudTable Alarm Rules	77
3.9.3 Viewing Cluster Monitoring Information	81
3.10 Audit Logs	83
3.10.1 Supported Key Audit Log Operations	83
3.10.2 Viewing Audit Logs	83
A Change History.....	86

1 Getting Started (Cluster Mode)

1.1 Creating a Cluster Quickly

This section describes how to create a cluster on the CloudTable management console.

Procedure

Step 1 Log in to the CloudTable management console at <https://console.huaweicloud.com/cloudtable/?locale=en-us>.

Step 2 Click **Purchase Cluster**. The **Purchase Cluster** page is displayed.

For example, the related parameter configurations are as follows:

- **Billing Mode:** Use the default value.
- **Region:** Use the default value.
- **AZ:** Use the default value.
- **Name:** Enter a cluster name, for example, **cloudtable-demo**.
- **I/O Type:** Use the default value.
- **VPC:** You can use the existing VPC. If you do not have a VPC, click **View VPC**. On the VPC management console that is displayed, create a VPC. Then you can select the newly created VPC from the VPC drop-down list on CloudTable. For details about how to create a VPC, see [Creating a VPC](#) in the *Virtual Private Cloud User Guide*.
- **Subnet:** When you create a VPC, a subnet is created by default. You can select a subnet name.
- **Security Group:** You can use the default value and ensure that the rules listed in [Figure 1-1](#) exist in the inbound rules of the selected security group.
 - Protocol & Port: All
 - Type: IPv4
 - **Source:** Local security group (that is, the security group to which the CloudTable cluster belongs)

Figure 1-1 Security group rule

Protocol & Port	Type	Source	Description	Operation
All	IPv4	Sys-default	-	Modify Replicate Delete

If the rule does not exist, add the preceding rules to the security group according to [Figure 1-2](#).


Figure 1-2 Adding an inbound rule

Protocol & Port	Source	Description	Operation
All ports	1-65535	Security...	Sys-default(2b486036-...)

If you are concerned that the default security group has network security risks, you can set as few open ports as possible in the security group. However, you need to ensure that the security group rules listed in [Table 1-1](#) exist in the selected security group.

Table 1-1 Customized 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	22	0.0.0.0/0	ECS remote login port (Linux)
	TCP	3389	0.0.0.0/0	ECS remote desktop service (Windows)
	TCP	16000	Security group of the CloudTable cluster	HMaster RPC port
	TCP	16010		HMaster web UI port
	TCP	16020		RegionServer RPC port
	TCP	16030		RegionServer web UI port
	TCP	2181		ZooKeeper client connection monitoring port
	TCP	2888		Follower connection monitoring port
	TCP	3888		ZooKeeper election port
	TCP	2000		HAgent access port
	TCP	8085		RestServer access port
	TCP	9600		RestServer access port
	TCP	4242		OpenTSDB access port

- **CloudTable Version:** Use the default value.
- **HBase Version:** Select 1.3.1 from the drop-down list.
- **Advanced Feature:** You can enable the following advanced features. If the OpenTSDB advanced feature is enabled, it will be charged. In this example, all advanced features are disabled.
 - OpenTSDB 2.3.0: It is an open source OpenTSDB time series database.
 - GeoMesa: It is a distributed, scalable, open source spatiotemporal database based on HBase.
- **IAM Authentication:** This function is disabled by default. If you need a security cluster with IAM authentication enabled, enable it.  shows that it is enabled.

Note: The IAM authentication function can be enabled or disabled only when you create a cluster. It cannot be enabled or disabled after the cluster has been created.

NOTE

The IAM authentication mode provides better security than the normal mode. Therefore, you are advised to enable IAM authentication for the CloudTable cluster and use IAM authentication in client or application code to connect to the cluster.

- **RS units:** Use the default value 2.

Step 3 Click **Next**.

Step 4 Confirm the details of the order and click **Submit**. The cluster creation task is submitted successfully.

Creating a cluster takes some time. The initial status of the cluster is **Creating**. After the cluster is created, the cluster status changes to **In service**.

----End

1.2 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 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 1-2 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
12005000210	J	Male	25	Dalian, Liaoning

Procedure

Step 1 Create a cluster named **cloudtable-demo**.

For details, see [Creating a Cluster Quickly](#).

Step 2 Prepare a Linux ECS.

Assume that the ECS name is **ecs_20170916**. For details, see [Preparing an ECS](#).

Step 3 Install a client and start the shell to access the CloudTable cluster.

For details about how to use the HBase shell to access the cluster, see [Using HBase Shell to Access a Cluster](#).

Step 4 Run the following commands on the HBase client to implement service A.

1. Create the **user_info** user information table and populate it with related data according to [Table 1-2](#).

```
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. Run the following command to delete the user information table.

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

Step 5 Delete the cluster.

For details, see [Deleting a Cluster](#).

----End

2 Permissions Management

2.1 Creating a User and Granting Permissions

This chapter describes [IAM](#) fine-grained permissions management for your CloudTable. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing CloudTable resources.
- Grant only the permissions required for users to perform a task.
- Entrust a HUAWEI CLOUD account or service to perform professional and efficient O&M on your CloudTable resources.

If your HUAWEI CLOUD account does not require individual IAM users, skip this chapter.

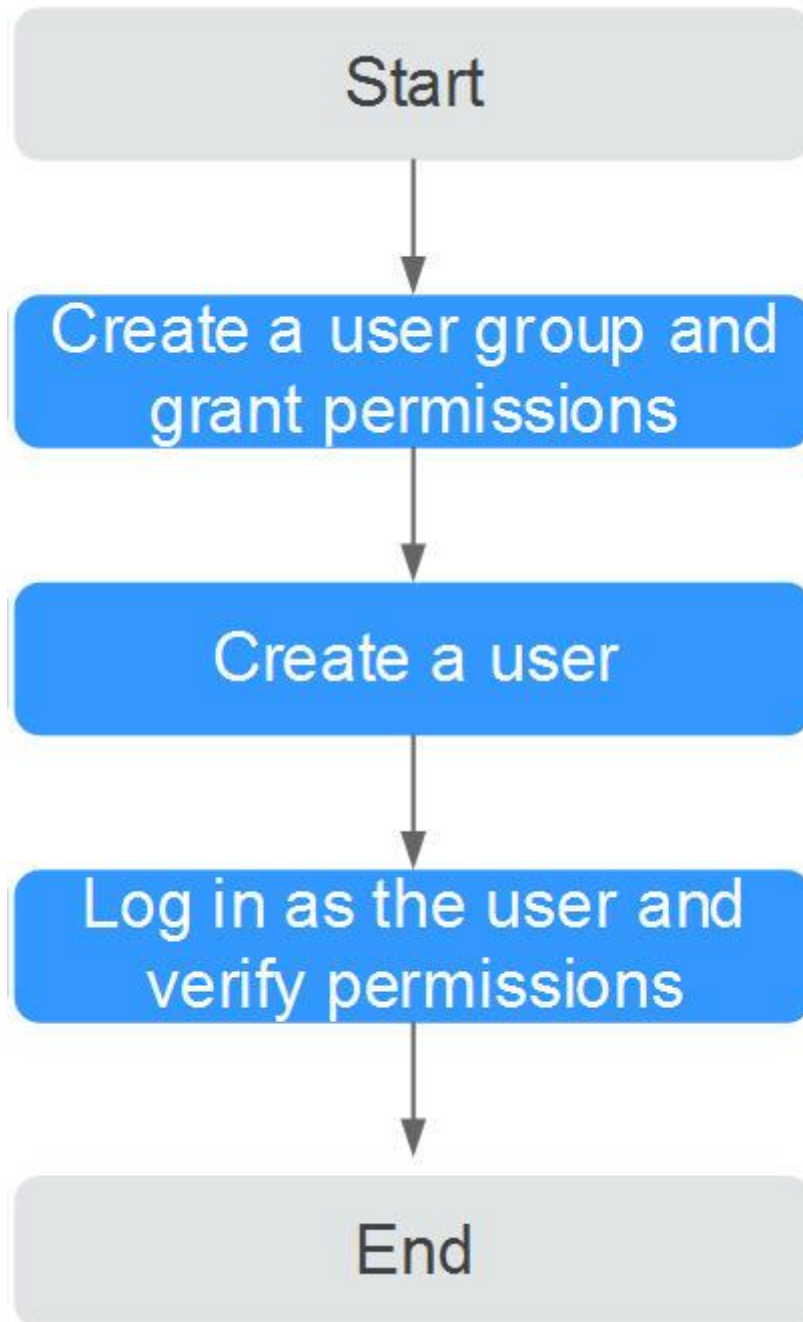
This section describes the procedure for granting permissions (see [Figure 2-1](#)).

Prerequisites

- CloudTable does not support fine-grained policies.
- Learn about the permissions (see [CloudTable System-Defined Permissions](#)) supported by CloudTable and choose policies or roles according to your requirements. For the system policies of other services, see [Permissions Policies](#).

Process Flow

Figure 2-1 Process for granting CloudTable permissions



1. [Create a user group and assign permissions](#) to it.
Log in to the IAM console, create a user group, and assign the **cloudtable Administrator** permission to it.
2. [Create a user and add it to a user group](#).
Create a user on the IAM console and add the user to the group created in Step 1.

3. [Log in](#) and verify permissions.

Log in to the console by using the user created, and verify that the user has the granted permissions.

- Choose **Service List > CloudTable Service**. Then click **Buy Cluster** on the CloudTable console. If the cluster can be created (for example, there is only the **cloudtable Administrator** permission), the **cloudtable Administrator** permission has already taken effect
- Choose any other service in the **Service List** (for example, there is only the **cloudtable Administrator** policy). If a message appears indicating insufficient permissions to access the service, the **cloudtable Administrator** policy has already taken effect.

3 Cluster Mode

3.1 Introduction to the Cluster Mode

The CloudTable cluster mode provides a distributed, scalable, and fully managed NoSQL data storage system based on Apache HBase. It provides strong consistency and single-digit millisecond latency so it is optimal for storage and queries of massive amounts of structured and semi-structured data. In addition, the CloudTable cluster mode integrates the following functions:

- **OpenTSDB 2.3.0:** It is an open source OpenTSDB time series database.
- **GeoMesa:** It is a distributed, scalable, open source spatiotemporal database based on HBase.

In CloudTable cluster mode, you need to create a dedicated cluster and can use it on demand. Dedicated clusters are physically isolated and are not affected by other services, facilitating user management. The CloudTable cluster mode is applicable to users with high service throughput and low delay requirements.

After a CloudTable cluster is created, you can access the CloudTable cluster using the following methods.

- Using a client to access the cluster
For details, see the following sections:
 - **Using HBase:** This section describes how to use the HBase shell to access a cluster.
 - **Using OpenTSDB:** This section describes how to use the cURL or Postman or compile code to send HTTP requests for calling APIs to access OpenTSDB.
 - **Using GeoMesa:** This section describes how to use the GeoMesa command line tool to access the GeoMesa spatiotemporal database.
- Call native open source APIs to develop HBase, OpenTSDB, and GeoMesa applications, and access the cluster through a CloudTable link.
For details, see [CloudTable Service Developer Guide](#).

The CloudTable cluster mode has been put into commercial use. For details about CloudTable pricing, see [Pricing Details](#).

3.2 Managing Clusters

3.2.1 Creating a Cluster

You can centrally manage clusters with CloudTable. A cluster is necessary for using CloudTable. This section describes how to create a cluster on the CloudTable management console.

The pay-per-use billing mode is used by default when you create a cluster. In pay-per-use mode, computing units are charged by the duration you use them, with a billing cycle of one hour. With this mode, you can start or stop a cluster at any time and pay what you use. You can also customize a CloudTable cluster with specified computing capabilities and storage space to meet your business needs. Alternatively, you can perform operations in [Purchasing a Discount Package](#) first and create a cluster to use the discount package billing mode. In this mode, you need to prepay nodes on a yearly or monthly basis, which is more economical than the pay-per-use mode.

Creating a Cluster


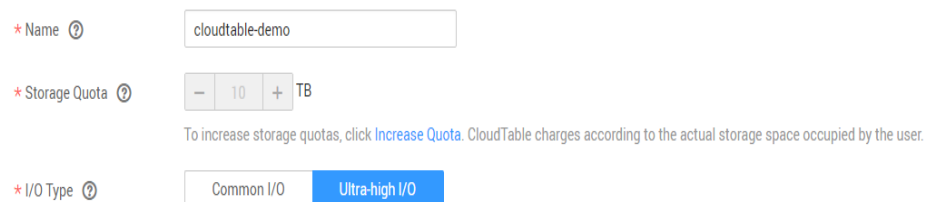
- Step 1** Log in to the CloudTable management console at <https://console.huaweicloud.com/cloudtable/?locale=en-us>.
- Step 2** Click  in the upper left corner to select a region.
- Step 3** Click **Purchase Cluster**. The **Purchase Cluster** page is displayed.
- Step 4** Configure basic cluster information by referring to the following table.

Table 3-1 Region

Parameter	Description
Region	Current working zone of the cluster. For more information about regions, see Regions and Endpoints .
AZ	Select the AZ associated with the cluster's region. For more information, see Regions and AZs .

Figure 3-1 Cluster configuration



The screenshot shows a configuration form with the following fields:

- Name:** A text input field containing "cloudtable-demo".
- Storage Quota:** A numeric input field with a minus sign on the left, "10" in the center, and a plus sign on the right, followed by "TB". Below it is a note: "To increase storage quotas, click [Increase Quota](#). CloudTable charges according to the actual storage space occupied by the user."
- I/O Type:** A radio button selection with two options: "Common I/O" (unselected) and "Ultra-high I/O" (selected and highlighted in blue).

Table 3-2 Cluster configuration

Parameter	Description
Name	Name of a cluster. A cluster name must start with a letter and contain 4 to 64 letters, digits, and hyphens (-). It cannot contain other special characters.
I/O Type	I/O type of the computing unit After the cluster is created, the I/O type cannot be changed. CloudTable supports the following I/O types: <ul style="list-style-type: none"> • Common I/O: uses Serial Advanced Technology Attachment (SATA) drives to store data. • Ultra-high I/O: uses solid state disk (SSD) drives to store data.

Figure 3-2 Network configuration




Table 3-3 Network configuration

Parameter	Description
VPC	A Virtual Private Cloud (VPC) is a secure, isolated, logical network environment. You can select an existing VPC or click View VPC to create a new one. For details about how to create a VPC, see Creating a VPC in the <i>Virtual Private Cloud User Guide</i> .
Subnet	A subnet provides dedicated network resources that are logically isolated from other networks, improving network security. A subnet is automatically created when a VPC is created. If you want to create a subnet, you can refer to Creating a Subnet for the VPC in the <i>Virtual Private Cloud User Guide</i> .
Security Group	A security group is used to control ECS access within a security group or between security groups by defining access rules. You can define different access control rules for a security group. These rules can specify which ECS ports or protocols are accessible and can be used to control inbound and outbound network traffic of VMs. After an ECS is added to the security group, it is protected by these access control rules. ECSs that do not belong to the security group cannot communicate with ECSs in the security group.

Parameter	Description
	<p>The underlying compute unit of a CloudTable cluster is ECS. In terms of security and stable service running, the ECSs must be added to the same security group in the same VPC. VPCs isolate networks and security groups specify which ports and protocols can be opened in the VPC.</p> <p>You can use an existing security group or click View Security Group to create a new one.</p> <p>For more information about security groups, see Security Group Overview in the <i>Virtual Private Cloud User Guide</i>.</p>

Table 3-4 Software configuration



Parameter	Description
CloudTable Version	CloudTable version
HBase Version	HBase version information.
Advanced Feature	<p>Advanced Feature: You can enable the following advanced features. If the OpenTSDB advanced feature is enabled, it will be charged.</p> <ul style="list-style-type: none"> OpenTSDB 2.3.0: It is an open source OpenTSDB time series database. GeoMesa: It is a distributed, scalable, open source spatiotemporal database based on HBase.
IAM Authentication	<p>If IAM authentication is required for a cluster, enable the IAM authentication function. Note: The IAM authentication function can be enabled or disabled only when you create a cluster. It cannot be enabled or disabled after the cluster has been created.</p> <ul style="list-style-type: none">  indicates that IAM authentication is disabled and the normal mode is used. This function is disabled by default.  indicates that IAM authentication is enabled. In this case, the client and applications must be configured with parameters such as the IAM username, accesskey, and secretkey to access the cluster. <p>NOTE The IAM authentication mode provides better security than the normal mode. Therefore, you are advised to enable IAM authentication for the CloudTable cluster and use IAM authentication in client or application code to connect to the cluster.</p>

Table 3-5 Computing unit

Parameter	Description
RS Units	Number of RegionServer units.
TSD Units	Number of OpenTSDB units. This parameter is displayed only when the

Parameter	Description
	OpenTSDB advanced feature is enabled.

Step 5 Click **Next**.

Step 6 Confirm the details of the order and click **Submit**. The cluster creation task is submitted successfully.

Step 7 Click **Back to Cluster List** to view the cluster status.

The cluster creation task takes some time. Please wait. The initial status of the cluster is **Creating**. After the cluster is created, the cluster status changes to **In service**.

----End

3.2.2 Purchasing a Discount Package

CloudTable also supports the monthly/yearly discount package. You can make a one-off payment according to the purchased service use duration. The service usage duration ranges from one month to one year. It is economical and recommended for long-term users. This section describes how to buy a CloudTable cluster in monthly/yearly discount package mode. After a discount package is purchased, the system does not automatically create a cluster. You need to create a cluster on the CloudTable management console.

Purchasing a Discount Package

Step 1 Log in to the [CloudTable management console](#).

Step 2 Click **Get Discounts**.

Step 3 On the **Buy CloudTable Service Discount Package** page, set the following parameters related to the discount package. The configuration fee is displayed at the bottom of the page. You can click **Pricing details** to view the charging details.

Figure 3-3 Discount package

Buy CloudTable Service Discount Package ⓘ

Region :
The purchased resource package applies to only the selected region. It cannot be shared across regions.

type: **Computing Unit**

Available Monthly: **750hours**

Duration: **1** | 2 | 3 | 4 | half year | 1 year

Specifications: ⓘ Within the validity period, fees are deducted from the purchased traffic package quota. Any fees beyond this are billed in pay-per-use mode.

Usage Duration: **1** | 2 | 3 | 4 | half year | 1 year

Purchase Quantity: - | 2 | +

Table 3-6 Parameters of the discount package

Parameter	Description	Example Value
Region	Select the region where the cluster runs. Resource packages in different regions are isolated. Select the region based on your requirements. For more information about regions, see Regions and Endpoints .	CN North-Beijing 1
Type	Computing unit. The computing units you purchase can be used as the computing units of RS and TSD.	Computing Unit
Available Monthly	750 hours per month	750 hours
Usage Duration	Select the usage duration.	-
Purchase Quantity	Number of computing units to be purchased. If you purchase the discount package for the first time, you are advised to purchase at least two computing units, because one CloudTable cluster requires at least two computing units. In addition, if the advanced feature OpenTSDB needs to be enabled during the CloudTable cluster creation, at least extra two computing units are required. If you have purchased the discount package before, select the number of computing units as required.	2

Step 4 Click **Next**. The **Confirm Order** page is displayed.

Step 5 Confirm the order information and click **Submit and Pay**.

Step 6 Choose either of the following payment modes:

- Choose **Balance** and click **Pay**.
- Choose **Online Payment** and click **Pay**.
- Choose **Request Online Contract and Payment** and click **Create Contract**. Enter the contract content and click **Create Formal Contract**.

After a discount package is purchased, the system does not automatically create a cluster. You need to create a cluster on the CloudTable management console. For details, see [Creating a Cluster](#).

---End

3.2.3 Introduction to Cluster Management

Log in to the CloudTable management console. In the left navigation pane, click **Cluster Mode**. On the displayed page, all CloudTable clusters are displayed. If there are a large number of clusters, you can turn pages to view clusters in any state.

Figure 3-4 Cluster list

Cluster Name	Cluster ...	Task Status	CloudTab...	Created	ZK Link	Billing ...	Operation
cloudtable-test		In service	v1.1.19	May 12, 2019 ...	cloudtable-e...	Pay-pe...	Restart View Metric More





Clusters are listed in chronological order by default, with the latest cluster displayed at the top. [Table 3-7](#) and [Table 3-8](#) provides parameters and icons in the cluster list.

Table 3-7 Cluster management parameters

Parameter	Description
Cluster Name	Name of a cluster. Set this parameter when creating a cluster.
Cluster Status	Cluster status, including Creating , In service , Sub-health , and Creation failed , and Frozen . For details, see Table 3-9 .
Task Status	Task status of a cluster For details, see Table 3-10 .
CloudTable Version	CloudTable version.
Created	Time when a cluster is created
ZK Link	ZooKeeper address
Billing Mode	Cluster billing mode
Operation	<ul style="list-style-type: none"> • Restart: Click Restart to restart a cluster. For details, see Restarting a Cluster. • Monitor: Click Monitor. The CloudTable Service Monitoring page is displayed.

Parameter	Description
	<ul style="list-style-type: none"> • More <ul style="list-style-type: none"> - Enable OpenTSDB: If you do not select OpenTSDB 2.3.0 when creating a cluster, Enable OpenTSDB will be displayed. You can click Enable OpenTSDB to enable it. For details, see Enabling OpenTSDB. - Expand capacity: Increase computing units in the cluster. For details, see Expanding Cluster Capacity. - Delete: You can click Delete to delete a cluster. For details, see Deleting a Cluster.

Table 3-8 Icon description

Icon	Description
	Enter a cluster name in the search box and click  to search for the cluster.
	Click  to manually refresh the cluster list.

Cluster Status

Table 3-9 Cluster status description

Status	Description
Creating	Indicates that a cluster is being created.
In service	If the cluster is successfully created and can provide services, the cluster status is In service .
Sub-health	If the cluster status cannot be monitored within the specified time, the cluster status changes to Sub-health . Manual intervention is required to recover a sub-healthy cluster. For example, you can restart the cluster to recover the cluster.
Creation failed	Indicates that a cluster fails to be created.
Frozen	<p>If the balance is insufficient for renewing a cluster, the cluster status is Frozen.</p> <p>To unfreeze the cluster, you need to recharge your account to ensure that the account balance is not 0. For details, see How Do I Renew the Service?. If you have purchased a discount package, you can unfreeze the cluster only after you recharge your account. Then you can purchase a discount package again. You cannot unfreeze a cluster by purchasing a discount package.</p> <p>NOTE</p> <p>A frozen cluster is unavailable and its all ECSs are shut down. After being unfrozen, the cluster recovers to the In service state. If you do not renew the</p>

Status	Description
	cluster before the freeze period ends, the cluster will be deleted.

Task Status

Table 3-10 Task status description

Status	Description
Deleting	Indicates that a cluster is being deleted.
Restarting	Indicates that a cluster is being restarted.
OpenTSDB enabling	Indicates that OpenTSDB is being enabled for the cluster.
Enabling OpenTSDB failed	Indicates that OpenTSDB fails to be enabled. You are advised to choose Resources > My Quota in the upper right corner of the management console to go to the Service Quota page. Check whether the remaining quota is sufficient. If the quota is insufficient, submit a work order to increase quotas. You can try again later or contact technical support.
Unit expanding	Indicates that a cluster is being expanded.
Unit capacity expansion failed	Indicates that the cluster fails to be expanded. You are advised to choose Resources > My Quota in the upper right corner of the management console to go to the Service Quota page. Check whether the remaining quota is sufficient. If the quota is insufficient, submit a work order to increase quotas. You can try again later or contact technical support.

3.2.4 Viewing Basic Cluster Information

You can monitor and manage the clusters you create. On the CloudTable management console, click **Cluster Mode**. In the cluster list, locate the cluster to be viewed and click the cluster name to access the basic information page. You can view the cluster information and network configurations.

[Table 3-11](#) and [Table 3-12](#) describe parameters about basic cluster information.

Table 3-11 Cluster information

Parameter	Description
Cluster Name	Name of a cluster. Set this parameter when creating a cluster.
Cluster ID	Cluster ID
Cluster Status	Cluster status information
Used Storage	Used storage capacity



Parameter	Description
Capacity (GB)	
RS Units	Number of RegionServer units
TSD Units	Number of OpenTSDB units
I/O Type	I/O type used by CloudTable
Billing Mode	Billing mode of the cluster
CloudTable Version	CloudTable version
Created	Time when a cluster is created
OpenTSDB Link	<p>OpenTSDB provides a web UI and you can access it using a URL.</p> <ul style="list-style-type: none"> • If OpenTSDB is enabled in the cluster, you can click  to copy the OpenTSDB link address. • After OpenTSDB is disabled in a cluster, you can search for the cluster in the cluster list on the Cluster Mode page. In the Operation column of the cluster, choose More > Enable OpenTSDB to enable OpenTSDB.
ZK Link	<p>ZooKeeper address.</p> <p>You can click  to copy the ZK link to the clipboard.</p>
IAM Authentication	Whether IAM authentication is enabled

Table 3-12 Network configuration

Parameter	Description
Region	Working area of the cluster. Set this parameter when creating a cluster.
AZ	AZ you select during cluster creation
VPC	<p>VPC you select during cluster creation</p> <p>A Virtual Private Cloud (VPC) is a secure, isolated, logical network environment.</p>
Subnet	<p>Subnet you select during cluster creation</p> <p>A subnet provides dedicated network resources that are logically isolated from other networks, improving network security.</p>
Security Group	Security group you select during cluster creation

3.2.5 Modifying HBase Parameters of the Cluster

Scenario

A CloudTable cluster has many HBase parameters related to read/write performance. You can modify the HBase parameters to tune cluster performance under different read/write request loads. You need to restart the cluster to make changes take effect.

Procedure

NOTE

- Restart the cluster after parameter configuration modification. Otherwise, service interruption may take place.
- Do not modify cluster parameters when the cluster is being restarted.

Step 1 Log in to the CloudTable management console.



Step 2 Click  in the upper left corner to select a region.


Step 3 In the left navigation pane, click **Cluster Mode**.

Step 4 Click the name of a cluster for which you want to modify HBase parameters to access the basic cluster information page.

Step 5 In the **Parameter Configuration** area, click the **Parameters** tab to modify HBase parameters.

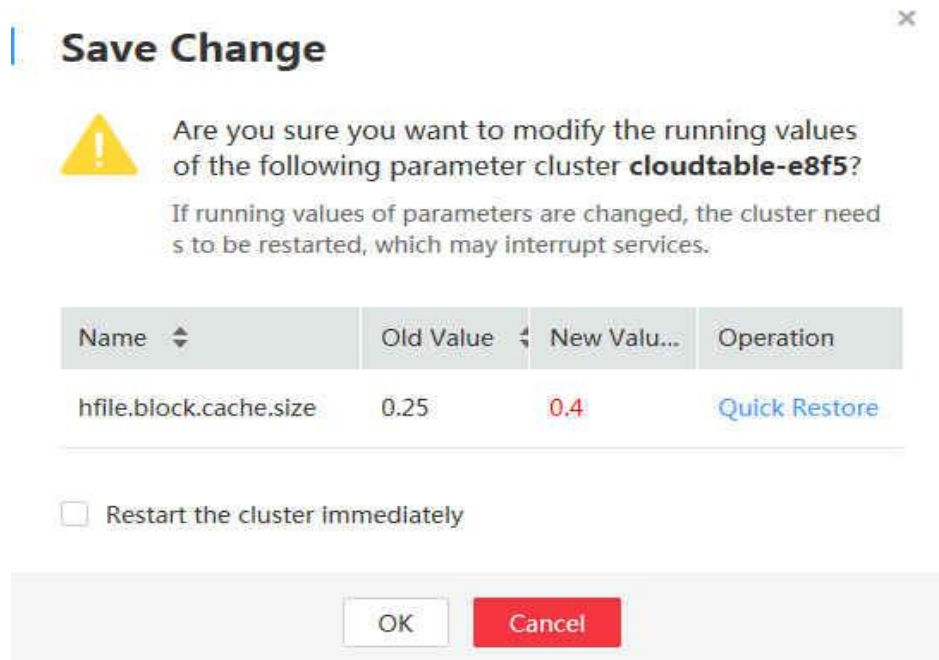
For details about the HBase parameters you can modify, see [HBase Parameters](#).


1. Select the target parameter and click  in the **Parameter Running Value** column.
2. Enter a new value in the text box, and then click . The parameter value has been successfully changed if the system prompts "The parameter changed to xx successfully. Save the modified value."

If you want to cancel it, click .

3. Click **Save Changes** in the upper left corner of the parameter list. The **Save Changes** dialog box is displayed.
4. In the **Save Changes** dialog box, verify whether the parameter settings are correct. Select **Restart the cluster immediately** and click **OK**.

Figure 3-5 Save changes



- After you select **Restart the cluster immediately** and click **OK**, the cluster restarts immediately. After the cluster restart is complete, the new parameter configurations take effect. Click  above the parameter list. The system prompts "The new value is applied" above the parameter list.
- If you do not select **Restart the cluster immediately** and click **OK** in the **Save Changes** dialog box, the system prompts "The new value is not applied" above the parameter list. In this case, you need to restart the cluster for the changes to take effect. For details about how to restart a cluster, see [Restarting a Cluster](#).
- If the new parameter value is incorrect, click **Quick Restore** to cancel the modification.

Step 6 After modifying parameters, you can click the **Change History** tab to view the change history.

On the **Change History** tab page, you can view the following information.

- **Name:** Name of the changed parameter.
- **Old Value:** Old parameter running value
- **New Value:** New parameter running value
- **Modification Time:** Time when you modify a parameter running value

----End

HBase Parameters

[Table 3-13](#) lists HBase parameters you can modify currently.

 NOTE

The sum of the values of **hbase.regionserver.global.memstore.size** and **hfile.block.cache.size** cannot exceed 0.8.

Table 3-13 HBase parameters

Parameter	Default Value	Value Range	Unit	Description
hbase.regionserver.thread.compaction.small	1	[1,20]	--	Indicates the number of HFile compaction threads. You can increase the parameter value in heavy-put-load scenarios.
hbase.regionserver.global.memstore.size	0.4	(0,0.8)	--	It is recommended that you set this parameter to "hbase.hregion.memstore.flush.size x Number of regions with active writes/RegionServer GC -Xmx)." The default value is 0.4 , indicating that 40% of RegionServer GC -Xmx is used.
hbase.hstore.blockingStoreFiles	60	[1,2147483647]	--	When the HFile number in the column cluster reaches this threshold, all operations in the region are blocked until the compaction is complete. You can increase the parameter value in heavy-put-load scenarios.
hbase.client.scanner.timeout.period	60000	[1,2147483647]	ms	A client and RegionServer parameter, indicating the scan lease period. It is recommended that you set this parameter to an integral multiple of 60000 ms, and increase the parameter value in heavy-read-load scenarios. Unit: ms
hfile.block.cache.size	0.2	(0,0.8)	--	Indicates the data cache percentage in the RegionServer GC -Xmx. You can increase the parameter value in

Parameter	Default Value	Value Range	Unit	Description
				heavy-read-load scenarios, in order to improve cache hit ratio and performance. The default value is 0.2, indicating that 20% of RegionServer GC -Xmx is used.
hbase.regionserver.handler.count	100	[1,300]	--	Indicates the number of RPC server instances on the RegionServer. The recommended value ranges from 100 to 300.
hbase.regionserver.metahandler.count	50	[1,100]	--	Indicates the number of program instances for processing prioritized requests. The recommended value ranges from 20 to 100.
hbase.hstore.flusher.count	2	[1,10]	--	Indicates the number of memstore flush threads. You can increase the parameter value in heavy-put-load scenarios.
hbase.ipc.server.callqueue.read.ratio	0.5	[0,1]	--	<p>When used under different load models, it controls the ratio between the numbers of read and write RPC queues. The value ranges from 0 to 1.0, and the default value is 0.5.</p> <p>0 indicates the total RPC queues of read and write operations.</p> <p>If the value is less than 0.5, it indicates that the read load is less than the write load.</p> <p>0.5 indicates that the read load equals to the write load.</p> <p>If the value is greater than 0.5, it indicates that the read load is greater than the write load.</p> <p>1.0 indicates that all</p>

Parameter	Default Value	Value Range	Unit	Description
				RPC queues except one are used for read operations.

3.2.6 Enabling OpenTSDB

If you do not select OpenTSDB when creating a CloudTable cluster, you can enable OpenTSDB after the cluster is created to provide the HTTP access function.

NOTE

Currently, OpenTSDB cannot be disabled once being enabled for a cluster.

Procedure

Step 1 Log in to the CloudTable management console.

Step 2 Click  in the upper left corner to select a region.

Step 3 In the left navigation pane, click **Cluster Mode**.

Step 4 In the cluster list, locate the row of the cluster for which OpenTSDB is to be enabled, and choose **More > Enable OpenTSDB**.

Step 5 In the dialog box that is displayed, set **TSD Units** and click **OK** to enable OpenTSDB.

In the **Task Status** column, the status of the OpenTSDB enabling task is displayed. The following provides task status details.

- **OpenTSDB enabling:** indicates that OpenTSDB is being enabled. Enabling OpenTSDB takes several minutes.
- If "--" is displayed following **OpenTSDB enabling**, OpenTSDB has been enabled successfully.
- **Enabling OpenTSDB failed:** indicates that OpenTSDB fails to be enabled. However, the failure does not affect the original status of the cluster, that is, **Cluster Status** is still **In service**.

You are advised to choose **Resources > My Quota** in the upper right corner of the management console to go to the **Service Quota** page. Check whether the remaining quota is sufficient. If the quota is insufficient, submit a service ticket to increase quotas. You can try again later or contact technical support.

----End

3.2.7 Restarting a Cluster

Procedure

Step 1 Log in to the CloudTable management console.

Step 2 Click  in the upper left corner to select a region.

Step 3 In the left navigation pane, click **Cluster Mode**.


Step 4 In the upper right corner of the cluster list, enter the name of a cluster in the search box and click .

Figure 3-6 Searching for a cluster by name



Step 5 In the **Operation** column of the cluster, click **Restart**.

Figure 3-7 Restarting a cluster



Step 6 In the dialog box that is displayed, select the check box and click **OK** to restart the cluster.

----End

3.2.8 Expanding Cluster Capacity

You can increase computing units to expand cluster capacity and improve system computing and storage capabilities. Increasing computing units is charged. You can dynamically increase the number of computing units based on site requirements or service conditions. The cluster automatically implements load balancing to ensure service continuity and smooth capacity expansion.

Before increasing TSD units, ensure that OpenTSDB has been enabled. If it is not enabled, follow instructions in [Enabling OpenTSDB](#) to enable it.

Precautions

- You can expand cluster capacity when a cluster is in the **In service** state and a task is not in any state or is in any of the following states: **Unit capacity expansion failed**, and **Enabling OpenTSDB failed**.
- The number of computing units to be added must be less than or equal to the remaining quotas. Otherwise, the system will prompt a message indicating that capacity expansion is not allowed.

Currently, each user's default computing unit quota is **32**. In the upper right of the management console, choose **Resources > My Quota**. On the displayed **Service Quota** page, view the used quota and total quota. If the quota is insufficient, click **Increase Quota** in the upper right of the page to increase computing unit quotas.

- During cluster capacity expansion, the system does not automatically restart the cluster, ensuring service continuity and smooth capacity expansion.
- If the capacity expansion fails, **Unit capacity expansion failed** is displayed in the **Task Status** column, and the cluster automatically rolls back to the state before capacity expansion. You can try again. If you have any questions, contact technical support.

Increasing Computing Units

Step 1 Log in to the CloudTable management console.

Step 2 Click  in the upper left corner to select a region.

Step 3 In the left navigation pane, click **Cluster Mode**.

Step 4 In the cluster list, locate the row that contains the target cluster, and choose **More > Expand Capacity**. **Expand Capacity** page is displayed.

Alternatively, you can click the name of the target cluster in the cluster list. On the displayed cluster details page, click **Expand Capacity** in the upper right part of the page.

Step 5 On the **Expand Capacity** page, set the following parameters and click **Next**.

Table 3-14 Capacity expansion parameters

Parameter	Description
Unit Type	Select the type of the unit to be added. The options are as follows: <ul style="list-style-type: none"> • RS Unit • TSD Unit: Before increasing TSD units, ensure that OpenTSDB has been enabled. If it is not enabled, follow instructions in Enabling OpenTSDB to enable it.
Expand to	Total number of units of the selected unit type after capacity expansion. A maximum of two computing units can be added at a time, and a maximum of 10 computing units can be added. A cluster supports a maximum of 100 computing units.

Step 6 Click **Submit** to start capacity expansion.

Step 7 Click **Back to Cluster List** to switch to the cluster list page.

If **Unit expanding** is displayed in the **Task Status** column, the cluster is being expanded. In this case, the cluster status is still **In service**.

After the capacity expansion is successful, click the cluster name in the cluster list. On the displayed cluster information page, you can view the number of computing units after capacity expansion.

----End

3.2.9 Deleting a Cluster

You can delete clusters you will no longer use. Deleting a CloudTable cluster will clear all resources and data related to the cluster. This operation cannot be undone. Exercise caution when deleting a cluster.

Procedure



- Step 1** Log in to the CloudTable management console.
- Step 2** Click  in the upper left corner to select a region.
- Step 3** In the left navigation pane, click **Cluster Mode**.
- Step 4** In the upper right corner of the cluster list, enter the name of a cluster in the search box and click .

Figure 3-8 Searching for a cluster



- Step 5** In the **Operation** column of the cluster, choose **More > Delete**.
 - Step 6** In the dialog box that is displayed, select the check box and click **OK** to delete the cluster.
- End

3.3 Preparing a Client Operating Environment

3.3.1 Preparing an ECS

If the client tool runs on Linux, you need to prepare a Linux ECS that is in the VPC as a CloudTable cluster and the Linux ECS serves as a client host. For example, if you use the following client tools to access a cluster, you are advised to use a Linux ECS.

- HBase shell
- GeoMesa command line tool

If the client tool runs on Windows, you need to prepare a Windows ECS that is in the VPC as a CloudTable cluster and the Windows ECS serves as a client host. Alternatively, you can refer to [Preparing a Local Windows Environment \(VPN Connection Mode\)](#) to prepare the client host. For example, if you use the following client tools to access a cluster, you are advised to use a Windows ECS.

- HBase web UI

You can select an ECS of different images and functions to meet your needs.

Preparing an ECS

For details about how to purchase a Linux or Windows ECS, see [Quickly Purchasing an ECS](#) in the *Elastic Cloud Server User Guide*.

The purchased ECS must meet the following requirements:

- The ECS must have the same region, AZ, VPC, and subnet as the CloudTable cluster.
For details about how to create a VPC, see [Creating a VPC](#) in the *Virtual Private Cloud User Guide*.
- The ECS must have the same security group as the CloudTable cluster.
For more information about security groups, see [Security Group Overview](#) in the *Virtual Private Cloud User Guide*.

Check whether the outbound rule of the security group contains the following rule:

- Protocol & Port: All
- Type: IPv4
- Destination: 0.0.0.0/0

If the preceding rules do not exist, add the following outbound rule according to [Figure 3-9](#).

- Protocol: TCP
- Port: 80
- Destination: Select **IP address** and set it to **0.0.0.0/0**.

Figure 3-9 Adding an outbound rule

Protocol & Port	Destination	Description	Operation
Custom TCP 80	IP address 0 . 0 . 0 . 0 / 0		Operation

- When purchasing an ECS, you need to set **EIP** to **Automatically assign**. Alternatively, you can bind an EIP to an ECS after the ECS is created.
- To access a Linux ECS, you need to log in to a Linux ECS using a password (SSH).
For details about how to log in to the ECS, see [Login Using an SSH Password](#) in the *Elastic Cloud Server User Guide*.
- Accessing a Windows ECS
For details about how to log in to the ECS, see [Login Using an MSTSC Password](#) in the *Elastic Cloud Server User Guide*.

Configuring the DNS Address and hosts File for the Linux ECS

You do not need to perform this operation when you install the HBase shell by deploying a client in one click. In other cases, perform the following steps to configure the DNS address and the **hosts** file for the Linux ECS:

Step 1 Configure `/etc/hosts`.

Configure the host IP address in the `/etc/hosts` file to accelerate the access to the CloudTable cluster.

To configure `/etc/hosts` for a Linux ECS, perform the following steps:

1. Log in to the Linux ECS as user **root**.
2. Run the **hostname** command to query the host name.

```
[root@euler ~]# hostname  
euler.novalocal
```

3. Run the **ifconfig** command to query the IP address of the local host.

Figure 3-10 Querying the IP address

```
[root@euler ~]# ifconfig  
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500  
    inet 192.168.0.58 netmask 255.255.255.0 broadcast 192.168.0.255  
    inet6 fe80::f816:3eff:fe8d:7acb prefixlen 64 scopeid 0x20<link>  
    ether fa:16:3e:8d:7a:cb txqueuelen 1000 (Ethernet)  
    RX packets 379443 bytes 522261185 (498.0 MiB)  
    RX errors 0 dropped 0 overruns 0 frame 0  
    TX packets 148302 bytes 10571485 (10.0 MiB)  
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0  
  
lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536  
    inet 127.0.0.1 netmask 255.0.0.0  
    inet6 ::1 prefixlen 128 scopeid 0x10<host>  
    loop txqueuelen 0 (Local Loopback)  
    RX packets 23819 bytes 16629808 (15.8 MiB)  
    RX errors 0 dropped 0 overruns 0 frame 0  
    TX packets 23819 bytes 16629808 (15.8 MiB)  
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0  
  
[root@euler ~]#
```

4. Run the **vi /etc/hosts** command to edit the file and add the host configuration.

```
192.168.0.58 euler.novalocal  
127.0.0.1 euler.novalocal
```

5. Press **Esc** and enter **:wq** to save the settings and exit.
6. Run the following command to check whether the IP address is successfully added:

```
cat /etc/hosts
```

7. Run the following command to check whether the host name can be resolved:

```
ping Host name
```

Step 2 Configure the DNS.

The DNS server is used to resolve the domain name in the CloudTable cluster link, for example, the ZooKeeper link. The IP addresses of the DNS server vary based on different regions.

- **CN North-Beijing1:** 100.125.1.250
- **CN East-Shanghai2:** 100.125.135.29,100.125.17.29
- **CN South-Guangzhou:** 100.125.1.250,100.125.21.250

To configure DNS for a Linux ECS, perform the following steps:

1. Log in to the Linux ECS as user **root**.
2. Run the **vi /etc/resolv.conf** command to edit the **/etc/resolv.conf** file and add the **nameserver** configuration.

Above the Internet IP address, add two lines to input the DNS server IP addresses in the following format.

```
nameserver 100.125.1.250
```

3. Press **Esc** and enter **:wq** to save the settings and exit.

4. Run the following command to check whether the IP address is successfully added:

```
cat /etc/resolv.conf
```

5. Run the following command to check whether an IP address can be resolved from the access domain name.

```
ping Access domain name
```

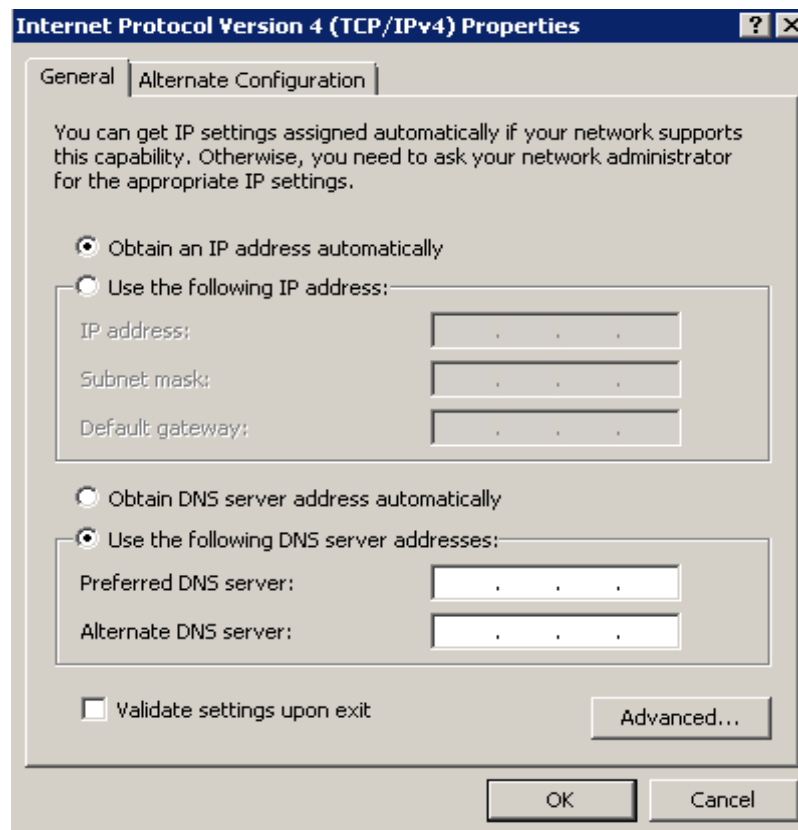
Access domain name indicates the domain name in the CloudTable cluster link, for example, the ZooKeeper link. On the CloudTable management console, click the cluster name to go to the basic cluster information page and obtain the link of CloudTable.

----End

Configuring DNS Address for a Windows ECS

1. Log in to the Windows ECS as user **Administrator**.
2. On the Windows ECS, choose **Start > Control Panel > Network and Internet > Network and Sharing Center**.
3. In the **Network and Sharing Center** window, click **Change adapter settings** on the right.
4. Right-click the network adapter name and choose **Properties** from the shortcut menu. The **Properties** dialog box is displayed.
5. Click **Internet Protocol Version 4 (TCP/IPv4)**. The **Internet Protocol Version 4 (TCP/IPv4) Properties** dialog box is displayed.

Figure 3-11 Internet protocol version 4 (TCP/IPv4) properties



6. In the dialog box shown in [Figure 3-11](#), select **Use the following DNS server addresses**, enter the IP address of the DNS server in the **Preferred DNS server** textbox, and click **OK**.

The DNS server is used to resolve the domain name in the CloudTable cluster link, for example, the ZooKeeper link. The IP addresses of the DNS server vary based on different regions.

- **CN North-Beijing1**: 100.125.1.250
- **CN East-Shanghai2**: 100.125.135.29,100.125.17.29
- **CN South-Guangzhou**: 100.125.1.250,100.125.21.250

7. Click **Start**, enter **cmd** in the search box, and press **Enter**. The CLI is displayed.
8. Run the following command on the CLI to check whether the DNS is successfully configured:

```
ping Access domain name
```

Access domain name indicates the domain name in the CloudTable cluster link, for example, the ZooKeeper link. On the CloudTable management console, click the cluster name to go to the basic cluster information page and obtain the link of CloudTable.

3.3.2 Preparing a Local Windows Environment (VPN Connection Mode)

If client tools run on Windows, you can use VPN on the local Windows host to connect to a CloudTable cluster through a Linux ECS.

Apply for a Linux ECS and bind it to an EIP. Install the IPsec VPN service on the Linux ECS so that it can be used as a VPN server. Use the local Windows environment as a client host and set up a VPN connection to connect to the VPN server (Linux ECS). After the connection is successful, you can connect to CloudTable clusters using the intranet link in the local environment. Before connecting to CloudTable, you need to install the JDK and client tools in the local Windows environment.

Advantages of the VPN connection mode:

1. IPsec is a secure VPN protocol and has two layers of authentication.
2. The function is stable. After the connection is successful, the internal IP address of the ECS in the CloudTable cluster can be properly accessed.

Preparing a Linux ECS and Installing and Configuring the VPN

- Step 1** Create and log in to a Linux ECS. The Linux ECS functions as a VPN server for installing the IPsec VPN service.

For details, see [Preparing an ECS](#). When creating an ECS, you need to bind an EIP. The ECS must have the same region, AZ, VPC, and subnet as the CloudTable cluster.

- Step 2** On the Linux ECS, download the installation script of the IPsec VPN service to the specified directory.

For CentOS, run the following command to download and execute the script:

```
wget https://git.io/vpnsetup-centos -O vpnsetup.sh
```

For Red Hat and SUSE, run the following command to download and execute the script:

```
wget https://git.io/vpnsetup -O vpnsetup.sh
```

Step 3 Set the pre-shared key (PSK), username, and password for the VPN service.

Run the **vi** command to open the **vpnsetup.sh** script, locate the parameters shown in the following figure, and modify their values as required.

Figure 3-12 Configuring VPN parameters

```
YOUR_IPSEC_PSK=' '  
YOUR_USERNAME=' '  
YOUR_PASSWORD=' '
```

You can also leave these parameters blank. The default username is **vpnuser**, and the password and PSK are generated randomly. After the installation is complete, a prompt is displayed, as shown in the following figure.

Figure 3-13 Prompt

```
Server IP: 11.11.11.77  
IPsec PSK: UAgRdsRoePyJhvNgq2GP  
Username: vpnuser  
Password: 9Kngrt7ASGtUqLh7
```

Step 4 Configure iptables firewall rules.

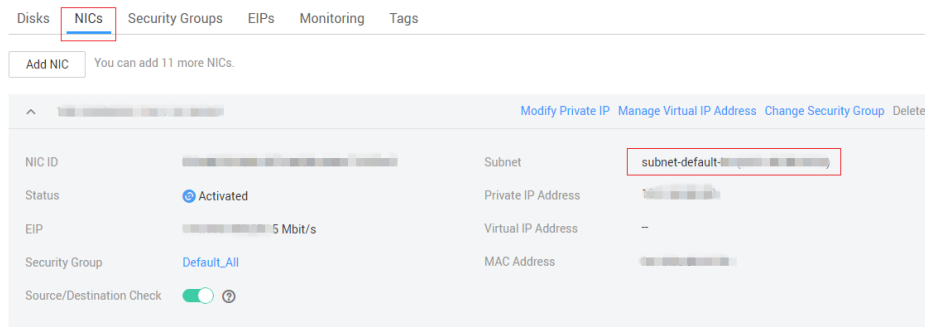
In the **vpnsetup.sh** script, locate the parameters shown in the following figure and set them to the iptables network segment that is the same as that of the Linux ECS subnet.

Figure 3-14 Configuring iptables rules

```
L2TP_NET=${VPN_L2TP_NET:-'192.168.42.0/24'}  
L2TP_LOCAL=${VPN_L2TP_LOCAL:-'192.168.42.1'}  
L2TP_POOL=${VPN_L2TP_POOL:-'192.168.42.10-192.168.42.250'}  
XAUTH_NET=${VPN_XAUTH_NET:-'192.168.43.0/24'}  
XAUTH_POOL=${VPN_XAUTH_POOL:-'192.168.43.10-192.168.43.250'}  
DNS_SRV1=${VPN_DNS_SRV1:-'8.8.8.8'}  
DNS_SRV2=${VPN_DNS_SRV2:-'8.8.4.4'}
```

The ECS subnet can be obtained from the NIC information in the ECS details.

Figure 3-15 NIC information



Step 5 Run the **vpnsetup.sh** script to install the IPsec VPN service and wait until the installation is complete.

```
sudo sh vpnsetup.sh
```

When the installation is complete, the **Server IP**, **IPsec PSK**, **Username**, and **Password** information is displayed. Record the information because you will use it when you set up a VPN connection.

After the installation is successful, you can run the **service ipsec status** command to check the service status. If the status is **active (running)**, the service is running properly.

----End

Setting Up a VPN Connection in the Local Windows Environment

The Windows 7 system is used as an example in the following operations.

Step 1 In the local Windows environment, you need to install the Java JDK first.

Visit the official website of Java to download the JDK 1.8.0 or a later version that matches the local Windows OS, and then install the JDK.

Step 2 Configure a security group rule to allow the local PC to connect to the VPN when accessing a CloudTable cluster.

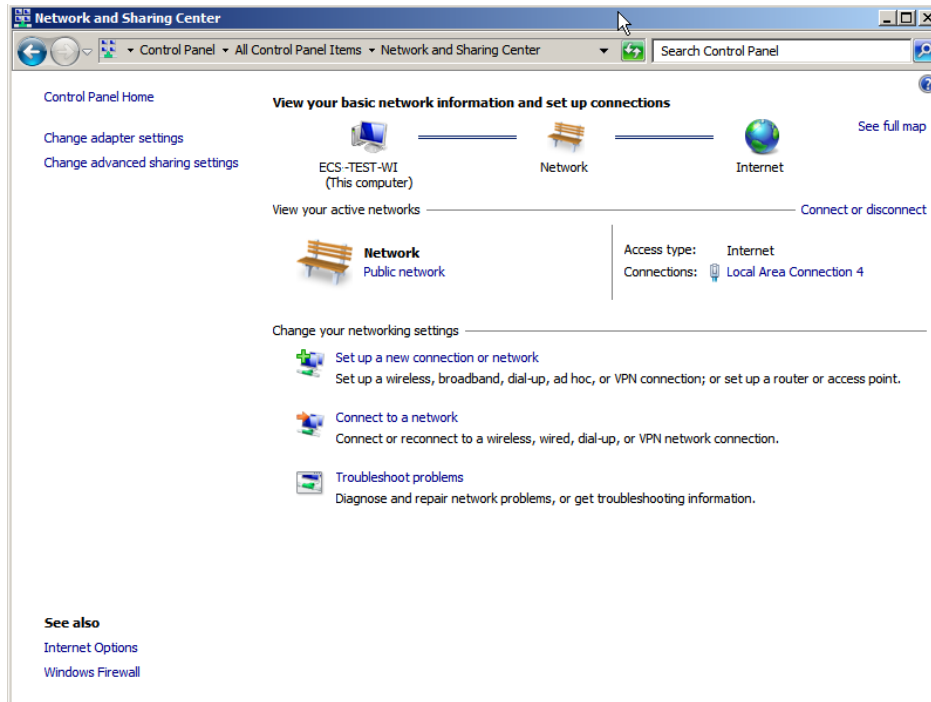
Log in to the CloudTable management console and choose **Cluster Mode**. Locate the required cluster in the cluster list, and click the cluster name to access the basic information page. In the **Network Configuration** area, click the security group name. On the displayed security group setting page, add the following two security group rules:

Table 3-15 Security group inbound rule

Direction	Protocol	Port/Range	Source/Security Group	Usage
Inbound	UDP	500	0.0.0.0/0	Allows the local PC to connect to the VPN when accessing a CloudTable cluster.
Inbound	UDP	4500	0.0.0.0/0	Allows the local PC to connect to the VPN when accessing a CloudTable cluster.

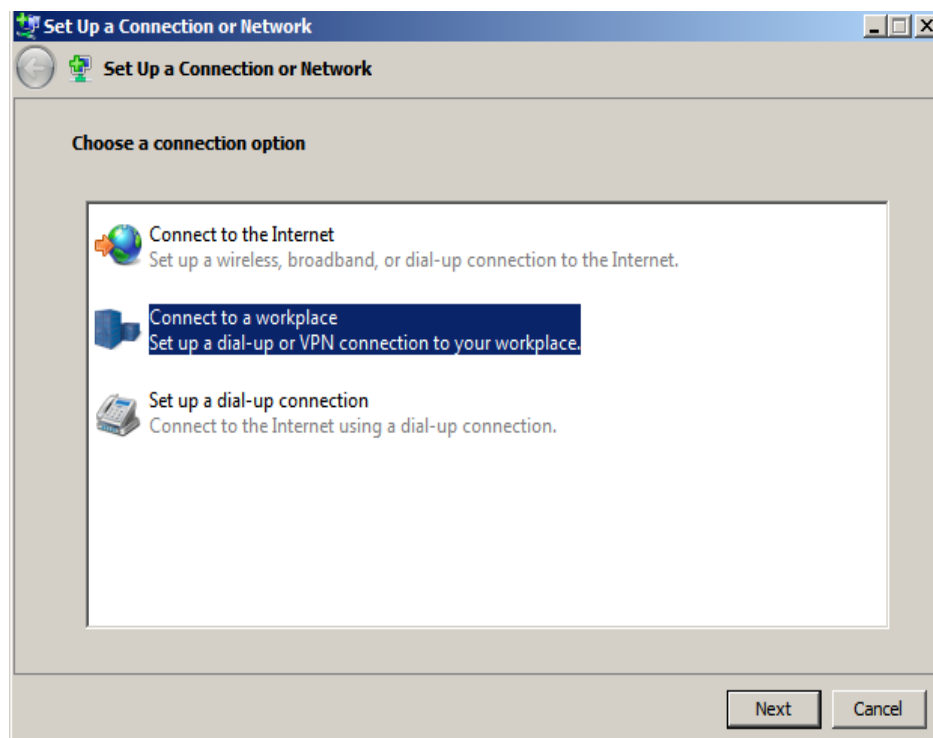
Step 3 In the local Windows environment, choose **Control Panel > Network & Internet > Network and Sharing Center** and click **Set up a new connection or network**.

Figure 3-16 Network and Sharing Center page



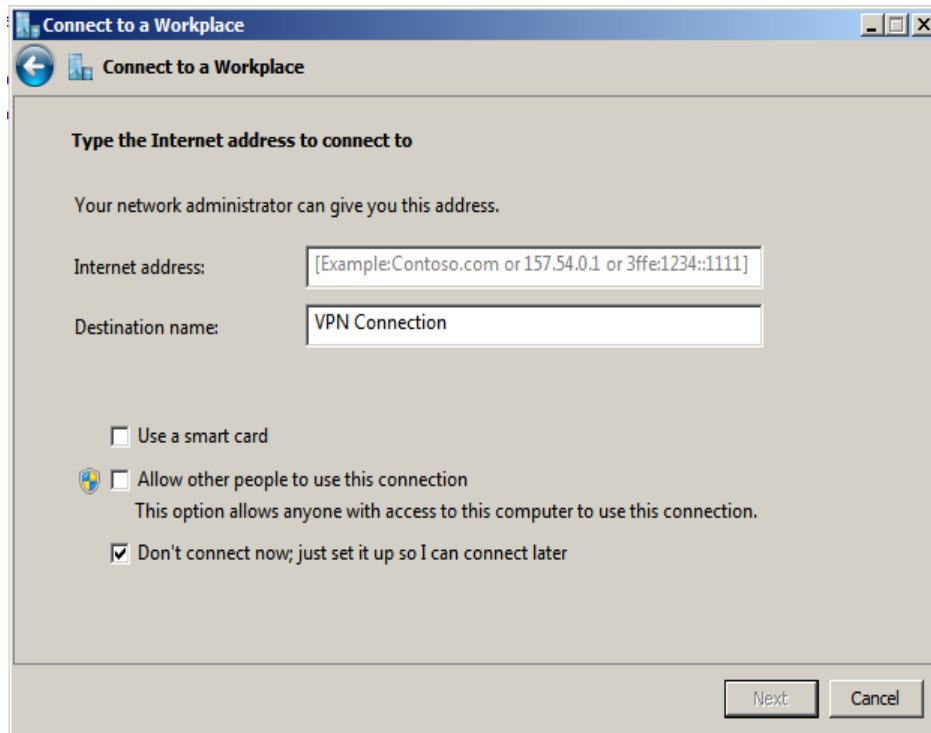
Step 4 In the displayed dialog box, select **Connect to a workplace**, click **Next**, and then select **Use my Internet connection (VPN)**.

Figure 3-17 Setting up a connection or network



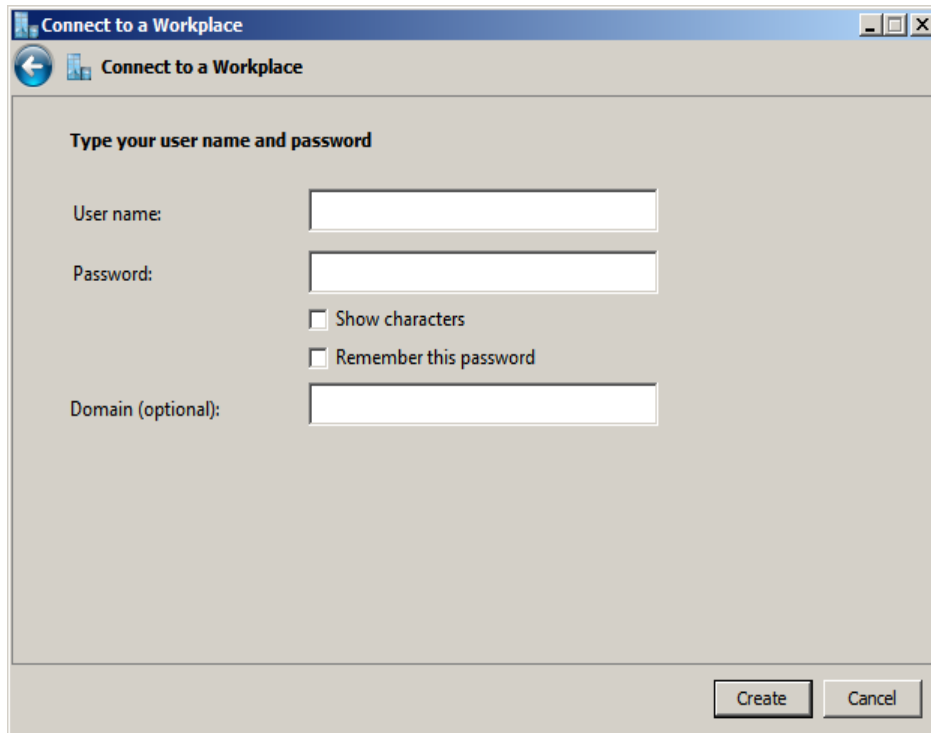
Step 5 Enter the EIP of the Linux ECS in the **Internet address** text box, select **Don't connect now; just set it up so I can connect later**, and click **Next**.

Figure 3-18 Setting a VPN connection



Step 6 Enter the VPN username and password set in [Step 3 in Preparing a Linux ECS and Installing and Configuring the VPN](#), and click **Create**. After the creation is successful, click **Close** to complete the creation of the VPN connection.

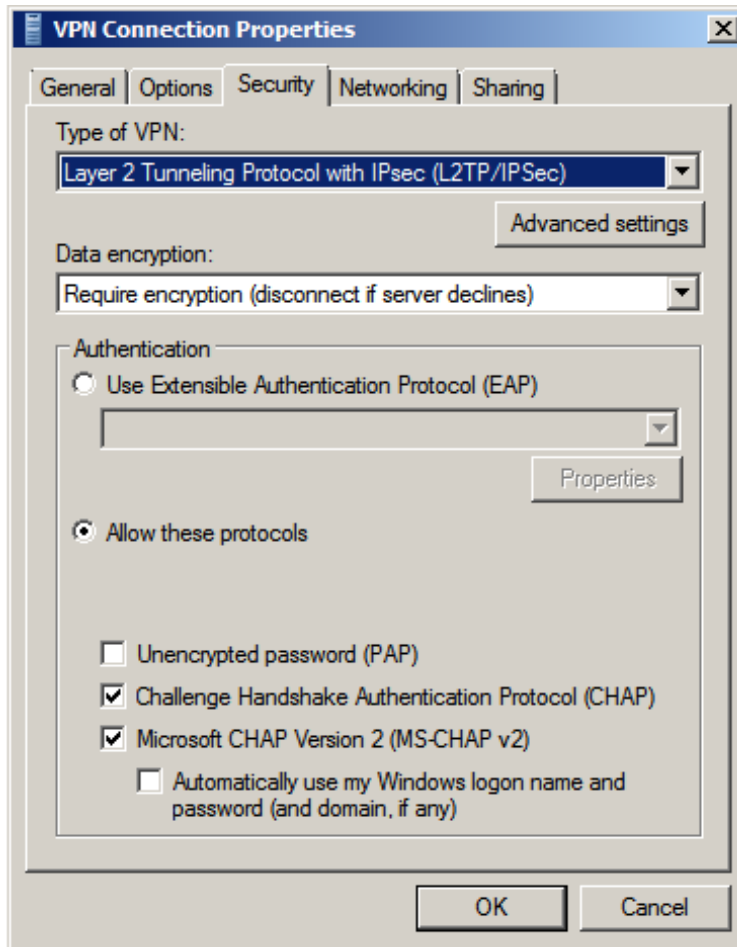
Figure 3-19 Setting the VPN username and password.



Step 7 Choose **Control Panel > Network & Internet > Connect to a network**. Right-click **VPN Connection** displayed in the lower right corner and choose **Properties** from the shortcut menu.

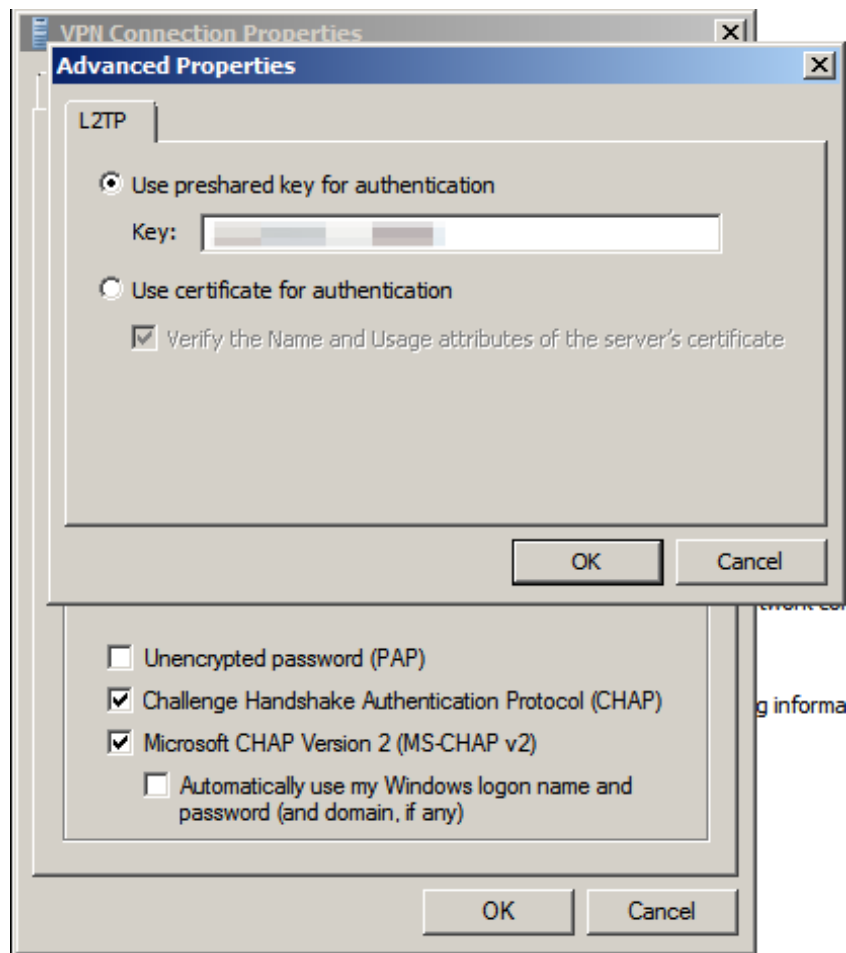
On the **Security** tab page of the **VPN Connection Properties** dialog box, set **Type of VPN** to **Layer 2 Tunneling Protocol with IPsec (L2TP/IPSec)**.

Figure 3-20 VPN connection properties



Click **Advanced settings**. In the **Advanced Properties** dialog box shown in the following figure, select the first option and enter the L2TP pre-shared key (PSK).


Figure 3-21 Advanced properties



Step 8 Log in to the local Windows OS as an administrator and click **Start**. Enter **cmd** in the search box and press **Enter** to open the command prompt window. Then, run the following statements to modify the registry. After the modification, restart the computer.

```
REG ADD HKLM\SYSTEM\CurrentControlSet\Services\PolicyAgent /v
AssumeUDPEncapsulationContextOnSendRule /t REG_DWORD /d 0x2 /f

REG ADD HKLM\SYSTEM\CurrentControlSet\Services\RasMan\Parameters /v ProhibitIpSec /t
REG_DWORD /d 0x0 /f
```

Step 9 Click the Internet access icon  in the lower right corner of the system, select the created VPN connection, and click **Connect** to connect to VPN. Wait until the VPN is successfully connected.

----End

Setting a Mapping Between the Intranet Domain Name and IP Address of CloudTable

Because the domain name in the link, for example, the ZK link of the CloudTable cluster is a HUAWEI CLOUD intranet domain name, the local PC cannot resolve it. Therefore, you need

to submit a service ticket to obtain the mapping between the domain name and the IP address, and configure it in the local **hosts** file.

- Step 1** Log in to the CloudTable management console, and choose **More > Service Tickets > Create Service Ticket** in the upper right corner. Enter the mapping between the intranet domain name of the CloudTable link and the IP address.
- Step 2** Add the mapping between the CloudTable intranet domain name and the IP address obtained in [Step 1](#) to the **hosts** file and save the file.

The following is an example of the mapping between the ZK link and the IP address of a CloudTable cluster. In the local Windows environment, open the **C:\Windows\System32\drivers\etc\hosts** file and add the following information:

```
192.168.0.239 cloudtable-demo-hmaster-2-1-Ie8OM83R.cloudtable.com
192.168.0.175 cloudtable-demo-regionserver-1-1-ZMD2WAG1.cloudtable.com
192.168.0.135 cloudtable-demo-regionserver-2-1-PRmohdbB.cloudtable.com
192.168.0.220 cloudtable-demo-hmaster-1-1-PJ0Qo8E0.cloudtable.com
192.168.0.147 cloudtable-demo-zk1-PJ0Qo8E0.cloudtable.com
192.168.0.231 cloudtable-demo-zk2-Ie8OM83R.cloudtable.com
192.168.0.232 cloudtable-demo-zk3-ZMD2WAG1.cloudtable.com
```

- Step 3** Click **Start** and enter **cmd** to open the CLI. Run the **ping** command to ping a CloudTable link. If it can be pinged, the CloudTable link can be accessed.
- End

3.4 Using HBase

3.4.1 Using HBase Shell to Access a Cluster

You can use the HBase shell to access a cluster by [deploying a client in one click](#) or [manually installing a client](#) on an ECS. [Deploying a Client in One Click](#) is recommended.

Deploying a Client in One Click

- Step 1** Prepare a Linux ECS.

Using a one-click client deployment tool. The recommended Linux ECS operating systems are EulerOS, CentOS, Ubuntu, and SUSE. For details, see [Preparing an ECS](#).

- Step 2** Download the one-click client deployment tool.


Use the SSH login tool (such as PuTTY) to remotely log in to the Linux ECS through the EIP and run the following command to obtain the one-click deployment tool.

```
curl -O -k "http://cloudtable-publish.obs.myhwclouds.com/quick_start_hbase_shell.sh"
```

- Step 3** Obtain a cluster access address.

Log in to the CloudTable management console and choose **Cluster Mode**. In the cluster list, locate the required cluster and obtain its ZK link in the **ZK Link** column. The parameter value is the cluster access address, as shown in [Figure 3-22](#).

Figure 3-22 ZK link



Step 4 Use the tool to deploy the client.

- Clusters with IAM authentication disabled

Replace **\$zookeeper_address** in the following command with the ZK link you obtained in [Step 3](#). Then, run the command on the CLI of the ECS to deploy the client in one click.

```
source quick_start_hbase_shell.sh $zookeeper_address
```

- Clusters with IAM authentication enabled

Specify account information for IAM authentication in the following command, and then run this command to deploy the client in one click.

```
source quick_start_hbase_shell.sh $zookeeper_address $user_name $access_key $secret_key
```

- **\$zookeeper_address**: Replace it with the ZK link obtained in [Step 3](#).
- **\$user_name**: Username used for creating a cluster. If the cluster is created using an account, set **\$user_name** to an account name. If the cluster is created using an IAM user under an account, set **\$user_name** to the *IAM username.account name*. An account and an IAM user are in a parent-child relationship.

NOTE

For a cluster with IAM authentication enabled, **\$user_name** must be set to the username used for creating the cluster and cannot be updated after the cluster is created. To update the username, submit a work order to the O&M personnel.

- **\$access_key** and **\$secret_key**: Access keys. Set **\$access_key** to the AK plaintext and **\$secret_key** to the SK plaintext. You can move your cursor over your account in the upper right corner of the management console and choose **My Credential**. Click **Access Keys** tab. On the **Access Keys** tab page, you can view the existing access keys or click **Add Access Key** to add an access key.

NOTE

The IAM authentication mode provides better security than the normal mode. Therefore, you are advised to enable IAM authentication for the CloudTable cluster and use IAM authentication in client or application code to connect to the cluster.

Step 5 Start the shell to access the cluster.

After you run the source command to automatically deploy the client, the HBase shell is automatically started. You can also run the **bin/hbase shell** command to start the HBase shell to access the cluster.

----End

Manually Installing a Client

Step 1 Prepare a Linux ECS.

For details, see the following content in [Preparing an ECS](#).

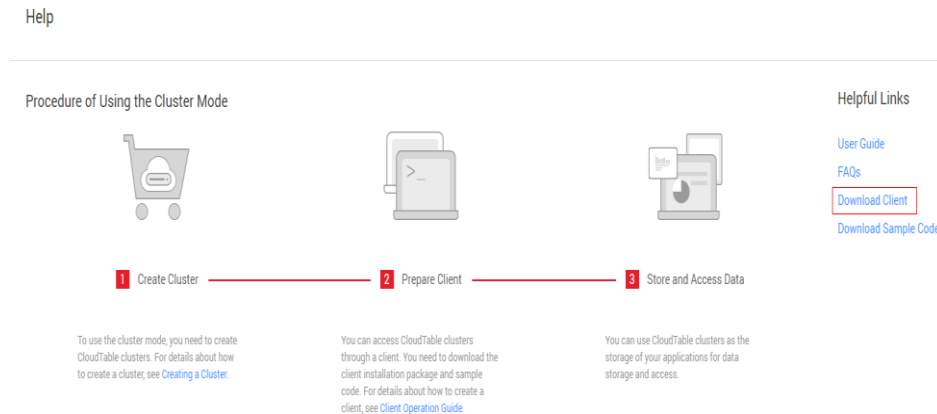
1. [Preparing an ECS](#)

2. Configuring the DNS Address and hosts File for the Linux ECS

Step 2 Download a client.

Log in to the CloudTable management console, click **Help** in the left navigation pane, and click **Download Client** on the right to download the client installation package.

Figure 3-23 Downloading the client



Step 3 Install the client.

1. Use a file transfer tool (such as WinSCP) to upload the client installation package to the Linux ECS.
2. 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 [Login Using an SSH Password](#) in the *Elastic Cloud Server User Guide*.

Run the following command to decompress the client installation package:

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

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


3. 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:

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

Log in to the CloudTable management console and choose **Cluster Mode**. In the cluster list, locate the required cluster and obtain its ZK link in the **ZK Link** column. See the following figure.

Figure 3-24 ZK link



- **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</va
lue>
</property>

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

Step 4 (Optional) Set security mode parameters.

Perform this step only when IAM authentication is enabled during cluster creation. Otherwise, skip this step.

1. Open the **hbase/conf/hbase-site.xml** file in the client installation directory.
2. Uncomment the security parameters except **hbase.zookeeper.quorum** and **mapreduce.cluster.local.dir**.

By default, the security parameters in the **hbase-site.xml** file are commented out. You need to uncomment the security parameters before configuring them.

3. Configure the username, accesskey, and secretkey.
 - **cloudtable.iam.username**: Username used for creating a cluster. If the cluster is created using an account, set **\$user_name** to an account name. If the cluster is created using an IAM user under an account, set **\$user_name** to the *IAM username.account name*. An account and an IAM user are in a parent-child relationship.
 - **cloudtable.iam.accesskey**: AK plaintext
 - **cloudtable.iam.secretkey**: SK plaintext

NOTICE

You can move your cursor over your account in the upper right corner and choose **My Credential**. Click the **Access Keys** tab and then the **Add Access Key** button to add an access key. After the access key is generated, you must download and save the key file that contains the AK plaintext and SK plaintext.

If you fail to query the AK plaintext and SK plaintext on the interface after the generation of the access key due to some reasons, for example, loss of the AK plaintext and SK plaintext, you can add another access key to make the previous one become invalid.

Retain the default values of other parameters.

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</va
lue>
</property>
```

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

<property>
  <name>cloudtable.iam.username</name>
  <value></value>
</property>
<property>
  <name>cloudtable.iam.accesskey</name>
  <value></value>
</property>
<property>
  <name>cloudtable.iam.secretkey</name>
  <value></value>
</property>
<property>
  <name>hbase.security.authorization</name>
  <value>true</value>
</property>
<property>
  <name>hbase.security.authentication</name>
  <value>digest</value>
</property>
<property>
  <name>hbase.rpc.protection</name>
  <value>authentication</value>
</property>
<property>
  <name>cloudtable.iam.service.password.encrypted</name>
  <value>false</value>
</property>
<property>
  <name>hbase.client.userprovider.class</name>
  <value>org.apache.hadoop.hbase.security.UserProviderExtend</value>
</property>
<property>
  <name>hbase.zookeeper.custom.auth.enabled</name>
  <value>true</value>
</property>
<property>
  <name>hbase.zookeeper.authProvider.scheme</name>
  <value>IAM AK SK</value>
</property>
```

Step 5 Start the shell to access the cluster.

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

---End

3.4.2 HBase Shell Commands

This section describes common HBase shell commands. For more HBase shell commands, visit <https://learnhbase.wordpress.com/2013/03/02/hbase-shell-commands/>.

1. Start the HBase shell.

Go to the HBase directory and run the following command to access the HBase shell:

```
./bin/hbase shell
```

2. Get help information.

After you run the **help** command in the HBase shell, all command information as well as common command instructions and use methods will be returned.

```
hbase(main):001:0> help
```

3. Create a table.

Run the **create** command to create a table. When creating a table, you must specify the table name and column family name.

```
hbase(main):007:0> create 'cloudtable','cf'
0 row(s) in 1.5530 seconds
=> Hbase::Table - cloudtable
```

4. Query a table:

Run the **list** command to query the specified table. The table name must be specified.

```
hbase(main):009:0> list 'cloudtable'
TABLE
cloudtable
1 row(s) in 0.0060 seconds
=> ["cloudtable"]
```

5. Insert a piece of record to a table.

Run the **put** command to insert a piece of record to the specified table. You need to specify the table name, primary key, customized column, and inserted value.

```
hbase(main):004:0> put 'cloudtable','row1','cf:a','value1'
0 row(s) in 0.2720 seconds
```

The following describes parameters in the command:

- **cloudtable**: table name
- **row1**: primary key
- **cf:a**: customized column
- **value1**: inserted value

6. Scan records

Run the **scan** command to scan a table. You need to specify the table name and you can scan a full table or specify a scanning range.

```
hbase(main):001:0> scan 'cloudtable'
ROW                                COLUMN+CELL
row1                                column=cf:a,
timestamp=1504866237162, value=value1
1 row(s) in 0.2420 seconds
```

7. Query a single record.

Run the **get** command to query a single record. The name and primary key of the queried table must be specified.

```
hbase(main):001:0> get 'cloudtable','row1'  
COLUMN                                CELL  
cf:a  
timestamp=1504866237162, value=value1  
1 row(s) in 0.2280 seconds
```

8. Disable a table.

Before modifying or deleting a table, you need to disable the table. Run the **disable** command to disable the table. When you perform operations on a disabled table, "ERROR" will be reported, indicating that the table is disabled.

```
hbase(main):002:0> disable 'cloudtable'  
0 row(s) in 2.3550 seconds
```

9. Enable a table.

If you want to use a table that has been disabled, run the **enable** command to enable the table.

```
hbase(main):004:0> enable 'cloudtable'  
0 row(s) in 1.2500 seconds
```

10. Delete a table.

Run the drop command to delete a table that you do not require any more. Before deleting a table, disable the table. Otherwise, "ERROR" will be displayed, indicating that the table is enabled. Deleting a table will cause data loss. Therefore, exercise caution when deleting a table.

```
hbase(main):007:0> disable 'cloudtable'  
0 row(s) in 2.2380 seconds  
  
hbase(main):008:0> drop 'cloudtable'  
0 row(s) in 1.2600 seconds
```

11. Exit the HBase shell.

Run the **quit** command to exit the HBase shell.

```
hbase(main):009:0> quit
```

3.4.3 Accessing the HBase Web UI

Scenario

The web UI displays the HBase cluster status, including cluster overview, RegionServer and master node information, snapshots, and running processes. You can learn about the overall HBase cluster status based on information displayed on the web UI.

Procedure

Step 1 Prepare a Windows ECS.

For details, see [Preparing an ECS](#) and [Configuring DNS Address for a Windows ECS in Preparing an ECS](#).

Step 2 In the browser address box, enter the following URL:


```
http://xxx-zk1.cloudtable.com:16010/master-status
```

In the preceding URL, **xxx-zk1.cloudtable.com** indicates the IP address of the active master node. The method of obtaining the master node address is the same as that of obtaining the ZooKeeper address. The zk1 or zk2 address in the ZooKeeper address list is the domain name of the master node.

You can obtain the ZooKeeper address (ZK link) from the basic cluster information page. For details, see [Viewing Basic Cluster Information](#).

Step 3 If you open the backup master page, you can click the link displayed in [Figure 3-25](#) to access the active master page.

Figure 3-25 Backup master page



- On the HBase web UI of the active master, the HBase summary is displayed on the homepage, including the following information:
 - The **Region Servers** page displays basic RegionServer information, as shown in [Figure 3-26](#).

Figure 3-26 Region Servers

Region Servers			
ServerName	Start time	Requests Per Second	Num. Regions
VM-3,21302,1415117599935	Wed Nov 05 00:13:19 CST 2014	0	1
VM-4,21302,1415117602775	Wed Nov 05 00:13:22 CST 2014	0	3
VM-5,21302,1415117607877	Wed Nov 05 00:13:27 CST 2014	0	0
Total:3		0	4

- The **Backup Master** page displays backup master information, as shown in [Figure 3-27](#).

Figure 3-27 Backup Masters

Backup Masters		
ServerName	Port	Start Time
VM-4	21300	Wed Nov 05 00:13:08 CST 2014
Total:1		

- The **Tables** page displays information about HBase tables, including **User Tables**, **Catalog Tables**, and **Snapshots**, as shown in [Figure 3-28](#).

Figure 3-28 Tables

Table Name	Online Regions	Description
t	1	't', {NAME => 't'}

- The **Tasks** page displays information about tasks running on HBase, such as the start time and status, as shown in [Figure 3-29](#).

Figure 3-29 Tasks

Start Time	Description	State	Status
Wed Nov 05 03:06:35 CST 2014	Closing region availabilityCheck_VM-5_1415127943_1415127994683.1d82d088bf4ba672ee2235fd98ae695.	COMPLETE (since 44sec ago)	Closed (since 44sec ago)
Wed Nov 05 00:20:13 CST 2014	RpcServer.reader=0,port=21300	WAITING (since 2mins, 36sec ago)	Waiting for a call (since 2mins, 36sec ago)
Wed Nov 05 00:19:49 CST 2014	RpcServer.reader=9,port=21300	WAITING (since 2mins, 53sec ago)	Waiting for a call (since 2mins, 53sec ago)
Wed Nov 05 00:19:17 CST 2014	RpcServer.reader=8,port=21300	WAITING (since 3mins, 8sec ago)	Waiting for a call (since 3mins, 8sec ago)
Wed Nov 05 00:15:10 CST 2014	RpcServer.reader=7,port=21300	WAITING (since 3mins, 46sec ago)	Waiting for a call (since 3mins, 46sec ago)
Wed Nov 05 00:14:52 CST 2014	RpcServer.reader=6,port=21300	WAITING (since 5mins, 2sec ago)	Waiting for a call (since 5mins, 2sec ago)
Wed Nov 05 00:14:36 CST 2014	RpcServer.reader=5,port=21300	WAITING (since 10sec ago)	Waiting for a call (since 10sec ago)
Wed Nov 05 00:14:01 CST 2014	RpcServer.reader=4,port=21300	WAITING (since 0sec ago)	Waiting for a call (since 0sec ago)

- The **Table Details** page displays overview of HBase storage tables, as shown in [Figure 3-30](#).

Figure 3-30 Table details

Table	Description
t	't', {NAME => 't', DATA_BLOCK_ENCODING => 'NONE', BLOOMFILTER => 'ROW', REPLICATION_SCOPE => '0', VERSIONS => '1', COMPRESSION => 'NONE', MIN_VERSIONS => '0', TTL => 'FOREVER', KEEP_DELETED_CELLS => 'false', BLOCKSIZE => '65536', IN_MEMORY => 'false', BLOCKCACHE => 'true'}

- The **Debug dump** page displays HBase debugging information, as shown in [Figure 3-31](#).

Figure 3-31 Debug dump

```
Master status for VM-5,21300,1415117586945 as of Wed Nov 05 03:13:34 CST 2014

Version Info:
=====
HBase V100R001C00
Subversion file:///home/Tomcat2/jenk_home/workspace/C00_FullBuild/hbase/source -r Unknown
Compiled by datasight on Mon Nov  3 21:24:59 IST 2014
Hadoop V100R001C00
Subversion git@rnd-git.huawei.com:datasight/hadoop2-4.git -r a39089a861837385c4d20a15fdaccb02442919a7
Compiled by datasight on 2014-11-03T15:47Z

Tasks:
=====
Task: RpcServer.reader=1,port=21300
Status: WAITING:Waiting for a call
Running for 10808s

Task: RpcServer.reader=2,port=21300
Status: WAITING:Waiting for a call
Running for 10805s

Task: RpcServer.reader=3,port=21300
Status: WAITING:Waiting for a call
Running for 10797s

Task: RpcServer.reader=4,port=21300
Status: WAITING:Waiting for a call
Running for 10772s

Task: RpcServer.reader=5,port=21300
Status: WAITING:Waiting for a call
Running for 10737s

Task: RpcServer.reader=6,port=21300
Status: WAITING:Waiting for a call
Running for 10721s

Task: RpcServer.reader=7,port=21300
Status: WAITING:Waiting for a call
Running for 10704s

Task: RpcServer.reader=8,port=21300
Status: WAITING:Waiting for a call
Running for 10456s

Task: RpcServer.reader=9,port=21300
Status: WAITING:Waiting for a call
Running for 10424s

Task: RpcServer.reader=0,port=21300
Status: WAITING:Waiting for a call
Running for 10400s
```

- The **HBase Configuration** page displays HBase configuration information, as shown in [Figure 3-32](#).

Figure 3-32 HBase Configuration

```
<?xml version="1.0" encoding="UTF-8"?>
- <configuration>
  - <property>
    <name>dfs.journalnode.rpc-address</name>
    <value>0.0.0.0:8485</value>
    <source>hdfs-default.xml</source>
  </property>
  - <property>
    <name>io.storefile.bloom.block.size</name>
    <value>131072</value>
    <source>hbase-default.xml</source>
  </property>
  - <property>
    <name>hbase.sessioncontrol.maxSessions</name>
    <value>65535</value>
    <source>hbase-site.xml</source>
  </property>
  - <property>
    <name>yarn.ipc.rpc.class</name>
    <value>org.apache.hadoop.yarn.ipc.HadoopYarnProtoRPC</value>
    <source>yarn-default.xml</source>
  </property>
  - <property>
    <name>mapreduce.job.maxtaskfailures.per.tracker</name>
    <value>3</value>
    <source>mapred-default.xml</source>
  </property>
  - <property>
    <name>hbase.rest.threads.min</name>
    <value>2</value>
    <source>hbase-default.xml</source>
  </property>
  - <property>
    <name>hbase.rs.cacheblocksonwrite</name>
    <value>>false</value>
    <source>hbase-default.xml</source>
  </property>
  - <property>
    <name>ha.health-monitor.connect-retry-interval.ms</name>
    <value>1000</value>
    <source>core-default.xml</source>
  </property>
  - <property>
    <name>yarn.resourcemanager.work-preserving-recovery.enabled</name>
    <value>>false</value>
    <source>yarn-default.xml</source>
  </property>
  - <property>
    <name>dfs.client.mmap.cache.size</name>
    <value>256</value>
    <source>hdfs-default.xml</source>
  </property>
  - <property>
```

----End

3.4.4 Authorizing Other Users to Access HBase in a Cluster with IAM Authentication Enabled

If you enable IAM authentication when creating a CloudTable cluster, other users have no permission to access HBase in the cluster unless they are authorized by you. If another user wants to access your CloudTable cluster with IAM authentication enabled, you need to run the **grant** command to grant permission to the user first. You can also run the **revoke** command to revoke the permission you grant to other users.

HBase Permission Control Commands

The **grant** command is used to grant specific permission to a specified user. Its syntax is as follows:

```
grant <user>, <permissions> [, <@namespace> [, <table> [, <column family> [, <column qualifier>]]]]
```

The **revoke** command is used to revoke the permission assigned to a specified user. Its syntax is as follows:

```
revoke <user> [, <@namespace> [, <table> [, <column family> [, <column qualifier>]]]]
```

user_permission is used to query permissions. Its syntax is as follows:

```
user_permission <table>
```

The **<permissions>** has five types of permissions:

- R - can read data within a given range.
- W - can write data within a given range.
- X - can have the execution permission within a given range.
- C - can create or delete tables within a given range (even those are created).
- A - can perform cluster operations, such as balancing loads between clusters or allocating regions within a given range.

For more information about HBase permission control, see [Access Control Labels \(ACLs\)](#) in the official HBase documentation.

Syntax examples:

```
--Grant the RWXC permission to userA.  
grant 'userA', 'RWXC'  
  
--View permissions of a specified table.  
user_permission 'table01'  
  
--Revoke the permission assigned to userA.  
revoke 'userA'
```

Authorizing Other Users to Access HBase in a Cluster with IAM Authentication Enabled

Assume that user A wants to access your cluster with IAM authentication enabled. You can perform the following steps to grant permission to user A and revoke the permission.

Step 1 Use the HBase shell to access the cluster. For details, see [Deploying a Client in One Click](#).

Step 2 In the HBase shell, run the following command to grant permission to user A:

, you can run the following command to authorize the user:

```
grant 'User A', 'RWXC'
```

Step 3 User A uses the HBase shell to access the cluster. For details, see [Deploying a Client in One Click](#).

When you run the following command to deploy the HBase shell in one click, replace the parameters in the command by referring to the parameter description below.

```
source quick start hbase shell.sh $zookeeper address $user name $access key  
$secret_key
```

- **\$zookeeper_address**: Enter the ZooKeeper link of the cluster with IAM authentication enabled you want to access.
- **\$user_name**: Username used for creating a cluster. If the cluster is created using an account, set **\$user_name** to an account name. If the cluster is created using an IAM user under an account, set **\$user_name** to the *IAM username.account name*. An account and an IAM user are in a parent-child relationship.
- **\$access_key** and **\$secret_key**: Enter the AK and SK of user A.

Step 4 After the HBase shell has been deployed in one click, user A can access the HBase shell to create tables and insert data.

For details about the HBase shell commands, see [HBase Shell Commands](#).

For example, run the following command to query all tables in HBase:

```
list
```

Step 5 If you want to forbid user A to access the cluster, you can revoke the permission.

You can run the following command in the HBase shell to revoke the permission:

```
revoke 'User A'
```

----End

Precautions

- If IAM authentication is not enabled in a CloudTable cluster, other users can access HBase in the cluster without authorization.
- If an IAM user under an account creates a CloudTable cluster with IAM authentication enabled, the account to which the IAM user belongs can access HBase in the cluster without authorization to read, write, and query HBase data. The account uses its AK and SK to connect to the ZK link of the IAM user's cluster.

3.5 Using OpenTSDB

3.5.1 OpenTSDB Overview

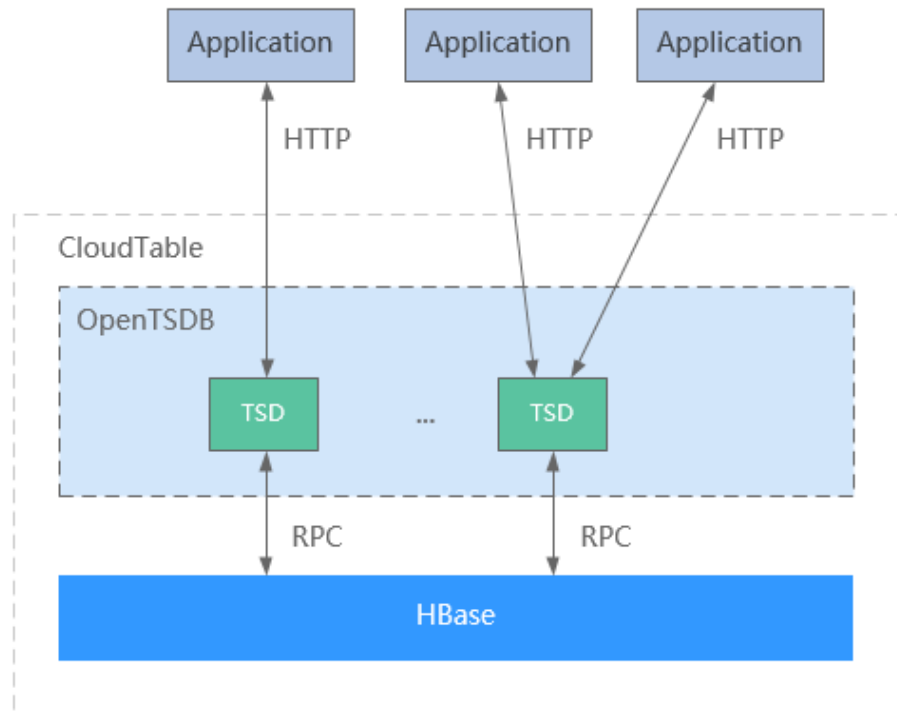
OpenTSDB is a distributed, scalable time series database based on HBase. It stores time series data. Time series data refers to the data collected at different time points. This type of data reflects the change status or degree of an object over time.

OpenTSDB Architecture

OpenTSDB consists of a Time Series Daemon (TSD) as well as a set of command line utilities. Interaction with OpenTSDB is primarily implemented by running one or more TSDs. Each TSD is independent. There is no master server and no shared state, so you can run as many TSDs as required to handle any load you throw at it. Each TSD uses HBase in a CloudTable cluster to store and retrieve time series data. The data schema is highly optimized for fast aggregations of similar time series to minimize storage space. TSD users never need to directly access the underlying storage. You can communicate with the TSD through an

HTTP API. All communications happen on the same port (the TSD figures out the protocol of the client by looking at the first few bytes it receives).

Figure 3-33 OpenTSDB architecture



Basic Concepts

- **data point:** A time series data point consists of a metric, a timestamp, a value, and a set of tags. The data point indicates the value of a metric at a specific time point.
- **metric:** Metrics include CPU usage, memory, and I/Os in system monitoring.
- **timestamp:** A UNIX timestamp (seconds or milliseconds since Epoch), that is, the time when the value is generated.
- **value:** The value of a metric is a JSON formatted event or a histogram/digest.
- **tag:** A tag is a key-value pair consisting of Tagk and Tagv. It describes the time series the point belongs to.

Tags allow you to separate similar data points from different sources or related entities, so you can easily graph them individually or in groups. One common use case for tags consists in annotating a data point with the name of the machine that produced it as well as name of the cluster or pool the machine belongs to. This allows you to easily make dashboards that show the state of your service on a per-server basis as well as dashboards that show an aggregated state across logical pools of servers.

Introduction to an OpenTSDB System Table

OpenTSDB stores time series data based on HBase. After OpenTSDB is enabled in a cluster, the system will create four HBase tables in the cluster. [Table 3-16](#) describes the OpenTSDB system tables.

NOTE

Do not modify the four HBase tables manually, because this may cause unavailable OpenTSDB.

Table 3-16 OpenTSDB system table

Table Name	Description
OPENTSDB.DAT A	It stores data points. All OpenTSDB data is stored in this table. OpenTSDB is partitioned based on salt. By default, 20 regions are supported. Currently, the number of regions cannot be configured.
OPENTSDB.UID	It stores unique identifier (UID) mappings. Each metric in a data point is mapped to a UID, and each tag in a data point is mapped to a UID. At the same time, each UID is reversely mapped to the metric or tag. These mappings are stored in this table.
OPENTSDB.TRE E	It stores metric structure information. This feature is disabled by default.
OPENTSDB.MET A	It stores time series indexes and metadata. This feature is disabled by default.

3.5.2 Connecting to OpenTSDB

If you enable OpenTSDB when creating a CloudTable cluster, you can access OpenTSDB through an OpenTSDB link on an ECS after the cluster is created to send HTTP requests to a resource path for data read and write.

If OpenTSDB is not enabled during cluster creation, you can enable it after the cluster is created by referring to [Enabling OpenTSDB](#).

Using HTTP APIs to Access OpenTSDB

OpenTSDB can be accessed only using HTTP APIs. For clusters with IAM authentication disabled, OpenTSDB uses HTTP for external access. For clusters with IAM enabled, OpenTSDB uses HTTPS for external access.

NOTE

HTTP has security risks and HTTPS is a secure protocol. You are advised to use HTTPS for connection.

OpenTSDB 2.3.0 is used in CloudTable and its APIs are consistent with the APIs of the open source OpenTSDB.

For more information about OpenTSDB APIs, see [OpenTSDB API Introduction](#) in the *CloudTable Service Developer Guide*.

Connecting to OpenTSDB

Step 1 Obtain an OpenTSDB link.

Log in to the CloudTable management console and choose **Cluster Mode** in the left navigation pane. Locate the cluster to be viewed in the cluster list and click the cluster name. On the basic information page that is displayed, obtain the OpenTSDB link (intranet).

Figure 3-34 OpenTSDB link



Step 2 Prepare a client operating environment.

The OpenTSDB link is an intranet address. You need to prepare an ECS that is in the same region, AZ, VPC, subnet, and security group as the CloudTable cluster to access OpenTSDB.

For details about how to create and log in to an ECS, see [Preparing an ECS](#).

Step 3 On the ECS, access OpenTSDB.

You can access OpenTSDB using either of the following methods:

- Use the [cURL](#) command line tool or a REST client such as [Postman](#) to send an HTTP request to access OpenTSDB.
For example, use the cURL command line tool to send the following request. If **200 OK** is returned, it indicates that OpenTSDB has been connected. **{OpenTSDB URL}** is the OpenTSDB link obtained in [Step 1](#).

```
curl -i -X GET http://{OpenTSDB URL}/api/version
```

The following provides examples of writing data and querying data:

a. Writing data:

```
curl -i -X POST -d '{"metric":"money", "timestamp":1524900283, "value":1, "tags":{"card":"card1"}}' http://{OpenTSDB URL}/api/put?sync -k -v
```

b. Querying data:

```
curl -i -X POST -d '{"start": 1524900283,"end": 1524900289,"queries":[{"aggregator": "sum","metric": "money"}]}' http://{OpenTSDB URL}/api/query -k -v
```

- Develop an application program to call HTTP APIs to access OpenTSDB for reading and writing data.

For details, see the following sections in the *CloudTable Service Developer Guide*:

- [Developing OpenTSDB Applications](#)
- [OpenTSDB API Introduction](#)

----End

3.6 Using GeoMesa

3.6.1 Using the GeoMesa Command Line Tools

You can use the GeoMesa shell or HBase shell to access a cluster by [deploying a client in one click](#) on an ECS.

Deploying a Client in One Click

Step 1 Prepare a Linux ECS.

Using a one-click client deployment tool. The recommended Linux ECS operating systems are EulerOS, CentOS, Ubuntu, and SUSE. For details, see [Preparing an ECS](#).

Step 2 Download the one-click client deployment tool.

Run the following command to obtain the one-click client deployment tool:

```
curl -O -k "http://cloudtable-publish.obs.myhwclouds.com/quick_start_hbase_shell.sh"
```

Step 3 Obtain a cluster access address.

Log in to the CloudTable management console and choose **Cluster Mode**. In the cluster list, locate the required cluster and obtain its ZK link in the **ZK Link** column. The parameter value is the cluster access address, as shown in [Figure 3-35](#).

Figure 3-35 ZK link

Cluster Na...	Cluster Sta...	Task Status	CloudTable...	Created	ZK Link (Intranet)	Operation
---------------	----------------	-------------	---------------	---------	--------------------	-----------

Step 4 Use the tool to deploy the client.

- Clusters with IAM authentication disabled

Replace **\$zookeeper_address** in the following command with the ZK link you obtained in [Step 3](#). Then, run the command on the CLI of the ECS to deploy the client in one click.

```
source quick_start_hbase_shell.sh $zookeeper_address
```

- Clusters with IAM authentication enabled

Specify account information for IAM authentication in the following command, and then run this command to deploy the client in one click.

```
source quick_start_hbase_shell.sh $zookeeper_address $user_name $access_key $secret_key
```

- **\$zookeeper_address**: Replace it with the ZK link obtained in [Step 3](#).
- **\$user_name**: Username used for creating a cluster. If the cluster is created using an account, set **\$user_name** to an account name. If the cluster is created using an IAM user under an account, set **\$user_name** to the *IAM username.account name*. An account and an IAM user are in a parent-child relationship.

NOTE

For a cluster with IAM authentication enabled, **\$user_name** must be set to the username used for creating the cluster and cannot be updated after the cluster is created. To update the username, submit a work order to the O&M personnel.

- **\$access_key** and **\$secret_key**: Access keys. Set **\$access_key** to the AK plaintext and **\$secret_key** to the SK plaintext. You can move your cursor over your account in the upper right corner of the management console and choose **My Credential**. Click **Access Keys** tab. On the **Access Keys** tab page, you can view the existing access keys or click **Add Access Key** to add an access key.

NOTE

The IAM authentication mode provides better security than the normal mode. Therefore, you are advised to enable IAM authentication for the CloudTable cluster and use IAM authentication in client or application code to connect to the cluster.

Step 5 Go to the directory where the GeoMesa client is installed.

After you run the source command to automatically deploy the client, the HBase shell is automatically started. Run the following command to exit the HBase shell client:

```
exit
```

Go to the directory where the GeoMesa client is installed.

```
cd geomesa-hbase_2.11-1.3.4
```

---End

3.6.2 GeoMesa Command Line

This section describes common GeoMesa commands. For more GeoMesa commands, visit <https://www.geomesa.org/documentation/user/accumulo/commandline.html>.

1. **View classpath.**

After you run the **classpath** command, all classpath information of the current command line tool will be returned.

```
bin/geomesa-hbase classpath
```

2. **Create a table.**

Run the **create-schema** command to create a table. When creating a table, you need to specify the directory name, table name, and table specifications at least.

```
bin/geomesa-hbase create-schema -c geomesa -f test -s  
Who:String,What:java.lang.Long,When>Date,*Where:Point:srid=4326,Why:String
```

3. **Describe a table.**

Run the **describe-schema** command to obtain table descriptions. When describing a table, you need to specify the directory name and table name.

```
bin/geomesa-hbase describe-schema -c geomesa -f test
```

4. **Import data in batches.**

Run the **ingest** command to import data in batches. When importing data, you need to specify the directory name, table name, table specifications, and the related data converter.

Data (license plate number, vehicle color, longitude, dimension, and time): **data.csv**. Save the data table to the **data** folder.

```
AAA,red,113.918417,22.505892,2017-04-09 18:03:46  
BBB,white,113.960719,22.556511,2017-04-24 07:38:47  
CCC,blue,114.088333,22.637222,2017-04-23 15:07:54  
DDD,yellow,114.195456,22.596103,2017-04-21 21:27:06  
EEE,black,113.897614,22.551331,2017-04-09 09:34:48
```

Table structure definition: **myschema.sft**. Save **myschema.sft** to the **conf** folder of the GeoMesa command line tool.

```
geomesa.sfts.cars = {  
  attributes = [  
    { name = "carid", type = "String", index = true }  
  ]  
}
```

```
{ name = "color", type = "String", index = false }
{ name = "time", type = "Date", index = false }
{ name = "geom", type = "Point", index = true, srid = 4326, default = true }
]
}
```

Converter definition: **myconvertor.convert**. Save **myconvertor.convert** to the **conf** folder of the GeoMesa command line tool.

```
geomesa.converters.cars= {
  type = "delimited-text",
  format = "CSV",
  id-field = "$fid",
  fields = [
    { name = "fid", transform = "concat($1,$5)" }
    { name = "carid", transform = "$1::string" }
    { name = "color", transform = "$2::string" }
    { name = "lon", transform = "$3::double" }
    { name = "lat", transform = "$4::double" }
    { name = "geom", transform = "point($lon,$lat)" }
    { name = "time", transform = "date('YYYY-MM-dd HH:mm:ss',$5)" }
  ]
}
```

Run the following command to import data:

```
bin/geomesa-hbase ingest -c geomesa -C conf/myconvertor.convert -s
conf/myschema.sft data/data.csv
```

For details about other parameters for importing data, visit

<https://www.geomesa.org/documentation/user/accumulo/examples.html#ingesting-data>.

5. Query explanations.

Run the **explain** command to obtain execution plan explanations of the specified query statement. You need to specify the directory name, table name, and query statement.

```
bin/geomesa-hbase explain -c geomesa -f cars -q "carid = 'BBB'"
```

6. Conduct statistical analysis.

Run the **stats-analyze** command to conduct statistical analysis on the data table. In addition, you can run the **stats-bounds**, **stats-count**, **stats-histogram**, and **stats-top-k** commands to collect more detailed statistics on the data table.

```
bin/geomesa-hbase stats-analyze -c geomesa -f cars
bin/geomesa-hbase stats-bounds -c geomesa -f cars
bin/geomesa-hbase stats-count -c geomesa -f cars
bin/geomesa-hbase stats-histogram -c geomesa -f cars
bin/geomesa-hbase stats-top-k -c geomesa -f cars
```

7. Export a feature.

Run the **export** command to export a feature. When exporting the feature, you must specify the directory name and table name. In addition, you can specify a query statement to export the feature.

```
bin/geomesa-hbase export -c geomesa -f cars -q "carid = 'BBB'"
```

8. Delete a feature.

Run the **delete-features** command to delete a feature. When deleting the feature, you must specify the directory name and table name. In addition, you can specify a query statement to delete the feature.

```
bin/geomesa-hbase delete-features -c geomesa -f cars -q "carid = 'BBB'"
```

9. Obtain the names of all tables in the directory.

Run the **get-type-names** command to obtain the names of tables in the specified directory.

```
bin/geomesa-hbase get-type-names -c geomesa
```

10. Delete a table.

Run the **remove-schema** command to delete a table. You need to specify the directory name and table name at least.

```
bin/geomesa-hbase remove-schema -c geomesa -f test
bin/geomesa-hbase remove-schema -c geomesa -f cars
```

11. Delete a catalog.

Run the **delete-catalog** command to delete the specified catalog.

```
bin/geomesa-hbase delete-catalog -c geomesa
```

3.7 HBase Elasticsearch Full-Text Search

3.7.1 Overview of Full-Text Search

HBase-Elasticsearch stores user source data in HBase and uses the Elasticsearch search engine of [Cloud Search Service](#) (CSS for short) to supplement full-text search based on key-value query capabilities. You can define which fields in HBase need full-text search based on service requirements. When you create an HBase table, a CSS cluster you specify will be automatically connected and an index is created in Elasticsearch. Index data is stored in Elasticsearch. In addition, the native APIs (Put and Scan) of HBase support the write and query of index data.

How to Use

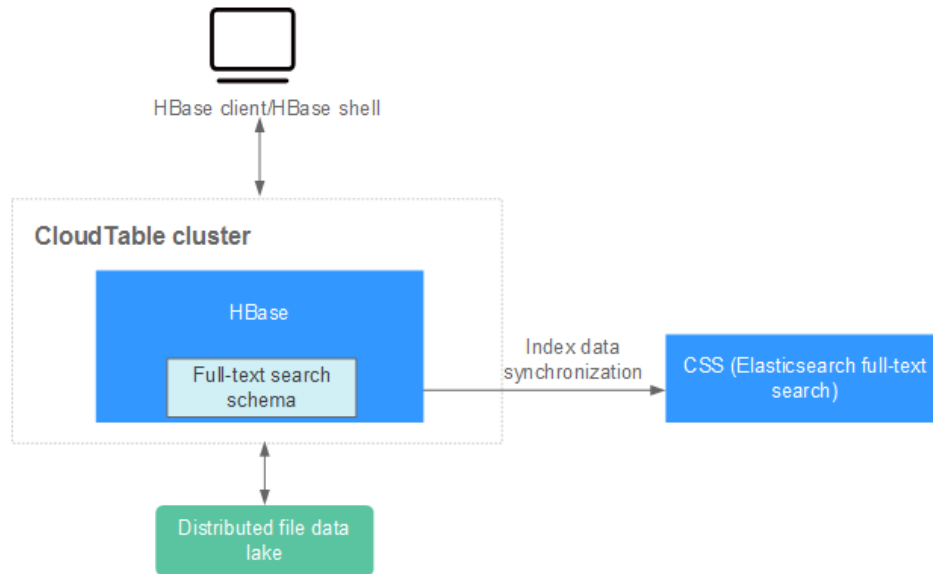
- [Using the HBase Shell for Full-Text Indexing](#)
- Java application development

You can develop an application to realize HBase Elasticsearch full-text search. For details, see [Developing HBase Elasticsearch Full-Text Search Applications](#) in the *CloudTable Service Developer Guide*.

Working Principles

As a big data storage service, CloudTable stores user data in bytes and provides efficient key-value random query capabilities. You can customize a schema to specify a data type (generally the text type) for some fields to extend the full-text search capability of CloudTable. CloudTable is suitable to be a primary storage system to store massive amounts of source data (any data type), because it separates computing from storage and features easy scale-out and cost-effectiveness of data storage. CSS (Elasticsearch) preserves lightweight index data to support keyword search. The following figure shows the working principles.

Figure 3-36 Working principles



If you enable a full-text index for some specified fields when creating an HBase table, HBase automatically synchronizes full-text index data to CSS when writing data. In addition, the native HBase data read API Scan supports common full-text search in terms of key-value read capability. To obtain complex high-level search capabilities, you can call Elasticsearch APIs and then CloudTable read APIs to complete service logic.

Application Scenarios

Massive amounts of user service data requires HBase to function as a big data online storage system to provide the most basic key-value query capabilities, featuring efficiency, high-concurrency, and low-latency. In addition, there are many types and quantities of fields in the data, that is, the corresponding services are diversified. For example, for a row of data in a table, some text fields need to use keywords for full-text search, some fields are secondary indexes, and some fields are applied to tag bitmap indexes. In this case, the Elasticsearch full-text search function needs to be enabled for CloudTable, while other service expansion capabilities are preserved. Example:

1. A search website stores massive amounts of search information, user environment information, and basic information in real time, extracts user information based on goods keywords, and resells the information to a third-party e-commerce platform.
2. An intelligent hospital's case system stores patients' medical treatment information, including the basic information, health status, doctor's occupational information, symptom description, diagnosis results, and medicine. A hospital information platform collects statistics on or searches for patients with historical medical treatment using keywords of the current social epidemics, prohibited drugs, or technical breakthroughs for tracking discharged patients or contacting patients to use new technologies for secondary diagnosis and other innovative services.
3. An intelligent public opinion governance system of governments stores massive amounts of data such as seditious speech, user information, and forwarding times of mainstream media platform users. It also searches for hot events in real time. If the event is a rumor, the system automatically reminds the user of the authenticity of the current event, the social impact data that the user publishes/forwards, relevant legal provisions, and similar

cases. The intelligent feedback mechanism is a deterrent to rumormongers and guides good public opinions.

HBase Elasticsearch Schema Definition

HBase uses metadata of a table to store the definition of the Elasticsearch schema.

Table 3-17 Schema definition

Field	Description	Mandatory
hbase.index.es.enabled	Whether to create a full-text index for the HBase table in Elasticsearch. The value true indicates that the full-text index is created. The default value is false .	Yes
hbase.index.es.endpoint	Access address of the CSS cluster (Elasticsearch engine), for example, ip1:port,ip2:port	Yes
hbase.index.es.indexname	Index name of the HBase table in Elasticsearch. The index name must be in lower case.	Yes
hbase.index.es.shards	Number of index shards in Elasticsearch. The default value is 5 . The value is an integer greater than or equal to 1.	No
hbase.index.es.replicas	Number of index replicas in Elasticsearch. The default value is 1 . The value is an integer greater than or equal to 0.	No
hbase.index.es.schema	Field mapping between HBase and Elasticsearch. The value is characters in JSON array format. Each element contains the following fields: <ul style="list-style-type: none"> name: Name of the field in Elasticsearch type: Type of the field in Elasticsearch hbaseQualifier: HBase qualifier of the data source analyzer: You can configure analyzer to specify an analyzer for fields of the text type. Typically, the ik_smart analyzer is used for Chinese text. The default value is Standard, supporting English text. <p>Example:</p> <pre>'[{"name":"contentCh","type":"text","hbaseQualifier":"cf1:contentCh","analyzer":"ik smart"}, {"name":"contentEng","type":"text","hbaseQualifier":"cf2:contentEng"}, {"name":"id","type":"long","hbaseQualifier":"cf1:id"}]'</pre>	Yes

The data types supported by HBase-Elasticsearch full-text search are {"text", "long", "integer", "short", "byte", "double", "float", "boolean"}, that is, the value type of **type** in the schema. **text** indicates the text type in Elasticsearch. Full-text search typically supports data of the text type and also supports accurate search of data of basic types.

3.7.2 Using the HBase Shell for Full-Text Indexing

This section describes how to use the HBase shell to create full-text indexes for HBase tables.

Prerequisites

The created CloudTable cluster (HBase), ECS instance (functioning as an HBase client), and CSS cluster (Elasticsearch engine) must have the same VPC, subnet, and security group to ensure network connectivity.

Full-text Search Example

Step 1 Start the HBase shell to access a CloudTable cluster.

For details about how to install and start the HBase shell, see [Using HBase Shell to Access a Cluster](#).

Step 2 Execute the following statement in the HBase shell to create an HBase table:

```
create 'hbase-es-table', {NAME => 'f', VERSIONS => 5}, SPLITS => ['10', '20'], METADATA => {'hbase.index.es.enabled' => 'true', 'hbase.index.es.endpoint' => '10.5.131.1:9200,10.5.131.2:9200', 'hbase.index.es.index.name' => 'product', 'hbase.index.es.schema' => '[{"name": "email", "type": "text", "hbaseQualifier": "f:email"}]'
```

For details about the schema definition of the **METADATA** field, see [HBase Elasticsearch Schema Definition](#). Replace **hbase.index.es.endpoint** in the preceding statement with the address to access the CSS cluster.

Step 3 In the HBase shell, run the following **put** commands to write three rows of data to the HBase table:

```
put 'hbase-es-table', '001rowkey', 'f:email', 'how many apples'
put 'hbase-es-table', '101rowkey', 'f:email', 'how much people'
put 'hbase-es-table', '201rowkey', 'f:email', 'many time people'
```

Step 4 Exit the HBase shell, and run the following **curl** command to call [Search APIs](#) of Elasticsearch and search for the keyword **how**:

```
curl -X GET "${ES_ClusterIP:Port}/product/search" -H 'Content-Type: application/json' -d '{ "storedfields" : ["rowkey"], "query" : { "term" : { "email" : "how" } } }'
```

Replace **ES_Cluster_IP:Port** in the preceding command with the address to access the CSS cluster, for example, **10.5.131.1:9200**.

Two documents (a document is a basic information unit for indexing and compiled in JSON) are hit in search result, and the rowkey field of the document is returned. The rowkey is the bridge of the mapping between HBase source data and Elasticsearch index data. The result is as follows:

```
{ "took":4, "timedout":false, "shards":{"total":5, "successful":5, "skipped":0, "failed":0 }, "hits":{"total":2, "maxscore":0.2876821, "hits":[{"index":"product", "type":"doc", "id":"GB087WYB7F1t0X-xu3ZX", "score":0.2876821, "fields":{"rowkey":["MDAxcM93a2V5"] } }, {"index":"product", "type":"doc", "id":"GR087WYB7F1t0X-xvHZ5", "_score":0.2876821, "fields":{"rowkey":["MTAxcM93a2V5"] } } ] } }
```

Step 5 Use the following website to perform decoding to obtain the rowkey of metadata in HBase:

<https://www.base64decode.org/>

The rowkey returned in [Step 4](#) is encoded by Base64.Encoder. You can obtain the rowkey in HBase using Base64.Decoder.

Step 6 Restart the HBase shell. In the HBase shell, run the following **get** command to obtain the data source:

```
get 'hbase-es-table','rowkey'
```

In Java application development, you can realize functions in [Step 3](#), [Step 4](#), and [Step 5](#) by invoking a function once. For details, see [Querying Data](#) in the *CloudTable Service Developer Guide*.

----End

3.8 Batch Data Import

3.8.1 Using CDM to Migrate Data to CloudTable

CloudTable uses CDM to migrate data from multiple data sources on the cloud and on-premises or third-party cloud to an HBase table of a CloudTable cluster. For details about the supported data sources, see [Data Sources Supported by CDM](#) in the *Cloud Data Migration User Guide*.

The procedure of using CDM to migrate data to CloudTable is as follows:

1. [Creating a CloudTable Cluster](#)
2. [Using CDM to Migrate Data to CloudTable](#)
3. [Viewing the Imported Data on CloudTable](#)

Creating a CloudTable Cluster

Step 1 Log in to the [CloudTable management console](#) to create a CloudTable cluster.

For details about how to create a CloudTable cluster, see [Creating a Cluster Quickly](#). If there is one available, you can skip this step.

For example, create a CloudTable cluster named **CloudTable-demo**.

----End

Using CDM to Migrate Data to CloudTable

Step 1 Log in to the [Cloud Data Migration management console](#) to create a CDM cluster.

The CDM cluster you create must have the same region, AZ, VPC, subnet, and security group as the CloudTable cluster.

For details about how to create a CDM cluster, see [Creating a Cluster](#) in the *Cloud Data Migration User Guide*.

For example, create a CDM cluster named **CDM-demo**.

Step 2 Create a source link in the new CDM cluster.

CDM enables you to migrate data from multiple data sources to CloudTable. Create a link as the source link based on the actual data source. The link is used for CDM to connect to the data source and read data from the data source.

For details about how to create a link, see [Creating a Link](#) in the *Cloud Data Migration User Guide*.

Step 3 Create a destination link in the new CDM cluster.

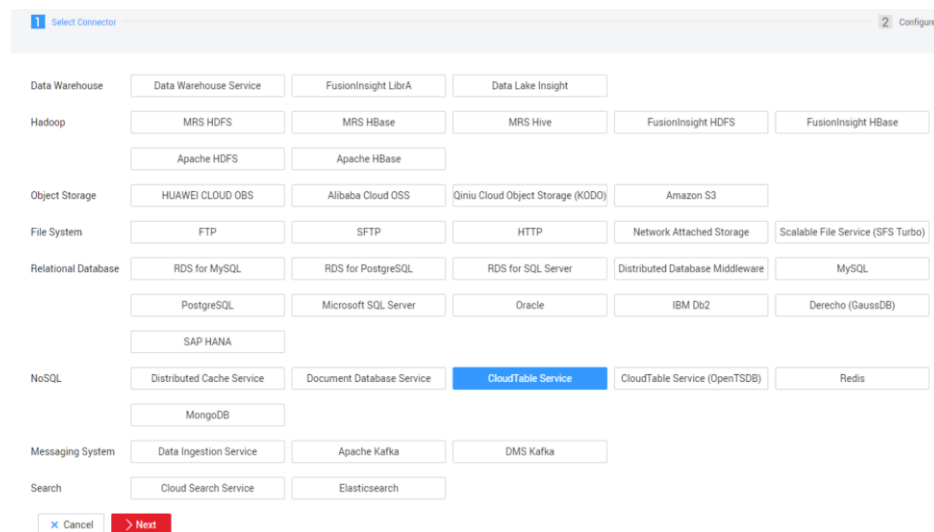
In the CDM cluster, create a link for which **Data Source Type** is set to **CloudTable** as a destination link. The link is used for CDM to connect to the CloudTable cluster, read data from the data source, and import data to the HBase table of the CloudTable cluster.

For details about how to create a link, see [Creating a Link](#) in the *Cloud Data Migration User Guide*.

The following is an example of destination link settings.

1. On the **Cluster Management** page, locate the newly created CDM cluster and click **Job Management** in the row of the CDM cluster. The job management page is displayed.
2. Click the **Link Management** tab and then click **Create Link**. On the **Select connector type** page that is displayed, select **CloudTable Service** and click **Next**.

Figure 3-37 Selecting a connector type



3. Configure link parameters.
 - **Name:** Enter a link name. You can create a link name based on the data source type for easy memorization. For example, enter **cloudtable_connect**.
 - **ZK Link:** Set this parameter to the ZK link of the CloudTable cluster. Log in to the CloudTable management console and choose **Cluster Mode**. In the cluster list, locate the required cluster and obtain its ZK link in the **ZK Link** column.
 - **IAM Authentication:** If you want to enable IAM authentication for the CloudTable cluster to be created, set this parameter to **Yes**. Otherwise, set this parameter to **No**.
 - **Run Mode:** Retain the default value.

Figure 3-38 Link parameter configuration

The screenshot shows a configuration form for a link. The fields are as follows:

- Name:** cloudtable_connector
- Connector:** HBase
- HBase Type:** CloudTable
- ZK Link:** cloudtable-demo-zk1.cloudt: Select
- IAM Authentication:** No
- Run Mode:** STANDALONE

At the bottom of the form, there are four buttons: Cancel, Previous, Test, and Save.

4. Click **Save**.

Step 4 In the CDM cluster, create a job on the **Table/File Migration** page.

For details about how to create a table/file migration job, see [Table/File Migration](#) in the *Cloud Data Migration User Guide*.

NOTE

When you create a table/file migration job, if the table or file of the source end does not have the corresponding HBase table in the destination CloudTable cluster, set a table name and field mapping in the job for the destination end (the field name can be copied from the source end). Therefore, when a job is running, the corresponding HBase table is automatically created at the destination end.

The following uses the SFTP data source on an ECS as an example to describe how to create a job.

1. On the **Cluster Management** page, locate the newly created CDM cluster and click **Job Management** in the row of the CDM cluster. The job management page is displayed.
2. Choose **Table/File Migration > Create Job**.
3. Set parameters as follows:

Figure 3-39 Configuring basic information

Job Configuration

* Job Name

Source Link Configuration

* Source Link Name

* Source Directory/File

* File Format

[Show Advanced Attributes](#)

Destination Link Configuration

* Destination Link Name

* Table Name

* Clear data before import Yes No

[Show Advanced Attributes](#)

4. Set the field mapping as follows:

Figure 3-40 Field mapping

Source Field			Destination Field			
Column ID	Example Value	Operation	Column Family	Column Name	Row Key	Operation
1	i	<input type="button" value="↺"/> <input type="button" value="🗑"/>	info	c1	<input type="checkbox"/>	<input type="button" value="🗑"/>
2	b	<input type="button" value="↺"/> <input type="button" value="🗑"/>	info	c2	<input type="checkbox"/>	<input type="button" value="🗑"/>
3	c	<input type="button" value="↺"/> <input type="button" value="🗑"/>	info	c3	<input type="checkbox"/>	<input type="button" value="🗑"/>

[Add Field+](#)

5. Configure the task as follows:

Figure 3-41 Configuring a task

Configure Task

Concurrent Extractors

Schedule Execution

[Show Advanced Attributes](#)

6. Click **Save and Run**.

----End

Viewing the Imported Data on CloudTable

Step 1 Prepare a Linux ECS.

Assume that the ECS name is `ecs_20170916`. For details, see [Preparing an ECS](#).

Step 2 Install a client and start the shell to access the CloudTable cluster.

For details about how to use the HBase shell to access the cluster, see [Using HBase Shell to Access a Cluster](#).

Step 3 On the CloudTable client, run a command to query the data migrated to CloudTable using CDM.

The following is an example of the command. Replace `table_name` with the name of the table specified when a job is created in the CDM cluster.

```
scan 'table_name'
```

----End

3.8.2 Using the Import Tool to Importing Data

Use the Import tool to import the **SequenceFile** data file of HBase to HBase of CloudTable. The Import tool is a tool contained in the installation package of the client tool.

NOTE

The **SequenceFile** file is a data file exported from HBase by the Export tool.

Step 1 Prepare a server where a CloudTable client is installed. For details about how to install a client, see [Using HBase Shell to Access a Cluster](#).

Step 2 Upload the **SequenceFile** file directory to the server where the CloudTable client is located. Assume that the directory is `/tmp/sequencefile`.

Step 3 The Import tool must be executed on the shell interface of the operating system of the client host. Go to the HBase directory on the client host and run the Import tool. The command format of the Import tool is as follows:

```
./bin/hbase org.apache.hadoop.hbase.mapreduce.Import <TableName> <InputPath>
```

TableName: Name of the table to which data is to be imported

InputPath: Directory of the **SequenceFile** file In the local file system, the directory must start with `file://`.

Example:

```
./bin/hbase org.apache.hadoop.hbase.mapreduce.Import t1 file:///tmp/sequencefile
```

----End

3.8.3 Using CopyTable to Import Data

CopyTable is a utility provided by HBase. It can copy part or of all of a table, either to the same cluster or another cluster. The target table must exist first. The CloudTable client tool includes CopyTable. After deploying the client tool, you can use CopyTable to import data to a CloudTable cluster.

Using CopyTable to Import Data

Step 1 Prepare a Linux ECS as the client host and deploy the CloudTable client tool on it.

For details, see [Using HBase Shell to Access a Cluster](#).

When deploying the client tool, set the ZK link to the ZK link of the CloudTable cluster where the source table resides.

Step 2 (Optional) If you want to copy a table to another cluster, obtain the ZK link of the target CloudTable cluster.

Log in to the CloudTable management console and choose **Cluster Mode**. In the cluster list, locate the required cluster and obtain its ZK link in the **ZK Link** column.

Step 3 Before using CopyTable to copy table data, ensure that the target table exists in the target CloudTable cluster. If the target table does not exist, create it first.

For details about how to create a table, see [Getting Started with HBase](#).

Step 4 On the client host, open the CLI, access the **hbase** directory in the installation directory of the client tool, and run the CopyTable command to import data to the CloudTable cluster.

The following is an example of the command. In this example, the data in the specified 1 hour in **TestTable** is copied to the target cluster.

```
cd ${Installation directory of the client tool}/hbase
./bin/hbase org.apache.hadoop.hbase.mapreduce.CopyTable --starttime=1265875194289
--endtime=1265878794289 --peer.adr=${ZK link of the target CloudTable cluster}:/hbase
--families=myOldCf:myNewCf,cf2,cf3 TestTable
```

----End

Overview of the CopyTable Command

The CopyTable command format is as follows:

```
CopyTable [general options] [--starttime=X] [--endtime=Y] [--new.name=NEW]
[--peer.adr=ADR] <tablename>
```

For details about the CopyTable command, see [CopyTable](#).

The following provides description about common options:

- **startrow**: the start row
- **stoprow**: the stop row
- **starttime**: beginning of the time range (unixtime in milliseconds). Without endtime means from the start time to forever.
- **endtime**: end of the time range. If no starttime is specified, ignore it.
- **versions**: number of cell versions to be copied
- **new.name**: name of a new table
- **peer.adr**: Address of the target cluster. The format is **hbase.zookeeper.quorum:hbase.zookeeper.client.port:zookeeper.znode.paren**. For the CloudTable clusters, the parameter value is **\${ZK link of the target CloudTable cluster}:/hbase**.
- **families**: List of column families to be copied. Multiple column families are separated by commas (,).

If you want to copy from **sourceCfName** to **destCfName**, specify **sourceCfName:destCfName**.

If the column family name needs to remain unchanged after copying, you only need to specify **cfName**.

- **all.cells**: Deletion markers and the deleted cells are also copied.

The parameter description is as follows:

tablename: name of the table to be copied

3.8.4 Sample Code for Copying a Table

You can develop an application and call HBase APIs to copy some or all data in a table either to the same cluster or another cluster. The target table must exist first.

Environment Preparation

- Step 1** If you want to copy a table to another cluster, obtain the ZK link of the target CloudTable cluster first.

Log in to the CloudTable management console and choose **Cluster Mode**. In the cluster list, locate the required cluster and obtain its ZK link in the **ZK Link** column.

- Step 2** Before copying table data, ensure that the target table exists in the target CloudTable cluster. If the target table does not exist, create it first.

For details about how to create a table, see [Getting Started with HBase](#).

- Step 3** Download the [sample code](#) and prepare the development environment by referring to [Preparing a Development Environment](#).

If you do not use the sample code project, copy the JAR file in the **cloudtable-example\lib** directory in the downloaded sample code to your project and add the JAR file to the dependency path in the project.

---End

Sample Code for Copying a Table

You can develop applications based on HBase APIs to copy table data to clusters based on service requirements. The following sample code is for reference only.

The following sample code is used to copy table data from one cluster to another cluster. IAM authentication is not enabled for both source and destination clusters.

```
package com.huawei.cloudtable.hbase.tool.client;

import java.io.IOException;
import java.util.ArrayList;
import java.util.Collections;
import java.util.List;
import java.util.concurrent.BlockingQueue;
import java.util.concurrent.LinkedBlockingQueue;
import java.util.concurrent.ThreadPoolExecutor;
import java.util.concurrent.TimeUnit;
import org.apache.commons.logging.Log;
```

```
import org.apache.commons.logging.LogFactory;
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.hbase.Cell;
import org.apache.hadoop.hbase.HBaseConfiguration;
import org.apache.hadoop.hbase.HRegionInfo;
import org.apache.hadoop.hbase.TableName;
import org.apache.hadoop.hbase.client.Admin;
import org.apache.hadoop.hbase.client.Connection;
import org.apache.hadoop.hbase.client.ConnectionFactory;
import org.apache.hadoop.hbase.client.Put;
import org.apache.hadoop.hbase.client.Result;
import org.apache.hadoop.hbase.client.ResultScanner;
import org.apache.hadoop.hbase.client.Scan;
import org.apache.hadoop.hbase.client.Table;

public class DataCoyer {
    public static final Log LOG = LogFactory.getLog(DataCoyer.class);
    private static int batchSize;
    private static int threads;

    public static void main(String[] args) throws IOException {
        if (args.length != 5) {
            System.err.println(
                "Command: ./hbase com.huawei.cloudtable.hbase.tool.client.DataCoyer [srcZK]
[dstZK] [tableName] [batchSize] [threads]");
            System.exit(-1);
        }
        batchSize = Integer.valueOf(args[3]);
        threads = Integer.valueOf(args[4]);
        copyData(args[0], args[1], args[2]);
    }

    public static void copyData(String srcZk, String dstZk, String tableName) throws
IOException {
        // Create a source HBase link.
        Configuration srcConf = HBaseConfiguration.create();
        srcConf.set("hbase.zookeeper.quorum", srcZk);
        //srcConf.set("zookeeper.znode.parent", "/hbase");
        Connection srcConnection = ConnectionFactory.createConnection(srcConf);

        // Create a target HBase link.
        Configuration dstConf = HBaseConfiguration.create();
        dstConf.set("hbase.zookeeper.quorum", dstZk);
        //dstConf.set("zookeeper.znode.parent", "/hbase 1.3.1");
        Connection dstConnection = ConnectionFactory.createConnection(dstConf);

        // Copy a table.
        copyDataTable(srcConnection, dstConnection, TableName.valueOf(tableName));
    }

    private static void copyDataTable(final Connection srcConnection, final Connection
dstConnection,
        final TableName tableName) throws IOException {
        try (Admin admin = srcConnection.getAdmin()) {
```

```
// Obtain the region information of the data table.
List<HRegionInfo> tableRegions = admin.getTableRegions(tableName);
Collections.shuffle(tableRegions);

// Create a thread pool and concurrently copy data.
BlockingQueue<Runnable> workQueue = new LinkedBlockingQueue<Runnable>(10000);
ThreadPoolExecutor executor = new ThreadPoolExecutor(threads, threads, 60,
TimeUnit.SECONDS,
    workQueue);

for (final HRegionInfo regionInfo : tableRegions) {
    executor.submit(new Runnable() {
        @Override
        public void run() {
            try (final Table srcTable = srcConnection.getTable(tableName);
                final Table dstTable = dstConnection.getTable(tableName)) {
                LOG.info("Start to copy region " + regionInfo.toString());
                Scan scan = new Scan();
                scan.setStartRow(regionInfo.getStartKey());
                scan.setStopRow(regionInfo.getEndKey());
                scan.setCaching(batchSize);
                copyOneRegion(srcTable, dstTable, scan, batchSize);
            } catch (IOException e) {
                LOG.error("CopyData failed .....", e);
                System.exit(-1);
            }
            LOG.info("Stop to copy region " + regionInfo.toString());
        }
    });
}

private static void copyOneRegion(Table srcTable, Table dstTable, Scan scan, int
batchSize)
    throws IOException {
    ResultScanner scanner = srcTable.getScanner(scan);
    Result result = null;
    List<Put> puts = new ArrayList<>();
    long counter = 0;
    while ((result = scanner.next()) != null) {
        Put put = new Put(result.getRow());
        for (Cell cell : result.rawCells()) {
            put.add(cell);
        }
        puts.add(put);

        counter += 1;
        if (puts.size() >= batchSize) {
            dstTable.put(puts);
            puts.clear();
        }
        if (counter % 10000 == 0) {
            LOG.info(srcTable + " has send: " + counter);
        }
    }
}
```

```

}
if (puts.size() != 0) {
    dstTable.put(puts);
}
}
}
}

```

3.9 Cluster Monitoring

3.9.1 Monitored CloudTable Metrics

Description

Monitoring is critical to ensure CloudTable reliability, availability, and performance. You can monitor the running status of CloudTable servers.

This section describes the metrics that can be monitored by Cloud Eye as well as their namespaces and dimensions. You can use the management console or [APIs](#) provided by Cloud Eye to query the metrics of the monitored objects and alarms generated for CloudTable.

Namespace

SYS.CloudTable

Monitored CloudTable Metrics

[Table 3-18](#) lists CloudTable metrics that can be monitored.

Table 3-18 Monitored CloudTable metrics

Metric ID	Metric	Meaning	Value Range	Monitored Object and Dimension	Monitoring Period (Raw Data)
cmdForNumberOfCpus	CPU Number	Number of CPUs of the instance	> 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
cmdProcessCPU	CPU Usage	CPU usage of the monitored object	0%~100%	Monitored object: CloudTable	1 minute

Metric ID	Metric	Meaning	Value Range	Monitored Object and Dimension	Monitoring Period (Raw Data)
		Unit: %		instance Dimension: instance_id	
cmdForTotalMemory	Total Memory	Total memory size of the monitored object Unit: byte	> 0 Byte	Monitored object: CloudTable instance Dimension: instance_id	1 minute
cmdProcessMem	Memory Usage	Memory usage of the monitored object Unit: %	0%~100%	Monitored object: CloudTable instance Dimension: instance_id	1 minute
cmdForIOWrite	Disks Write Rate	Volume of data written to the monitored object per second Unit: byte/s	>= 0 Byte/s	Monitored object: CloudTable instance Dimension: instance_id	1 minute
cmdForIORead	Disks Read Rate	Volume of data read from the monitored object per second Unit: byte/s	>= 0 Byte/s	Monitored object: CloudTable instance Dimension: instance_id	1 minute
cmdForUsedStorageSize	Used Storage Size	Used storage space in the cluster Unit: byte	>= 0 Byte	Monitored object: CloudTable instance Dimension: instance_id	1 minute
cmdForUsedStorageRate	Used Storage Rate	Ratio of the used storage space to the total storage space in the cluster	0%~100%	Monitored object: CloudTable instance	1 minute

Metric ID	Metric	Meaning	Value Range	Monitored Object and Dimension	Monitoring Period (Raw Data)
		Unit: %		Dimension: instance_id	
hm_regionservernum	Normal RegionServers	Number of normal RegionServers	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
hm_deadregionservernum	Faulty RegionServers	Number of faulty RegionServers	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
hm_ritCountOverThreshold	RIT Count Over Threshold	Region in transaction count over threshold	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
hm_ritCount	RIT Count	Region in transaction count	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_requests	Requests Per Second	Number of requests of a RegionServer per second Unit: Request/s	>= 0 request/s	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_regions	Regions	Number of regions of a RegionServer	>= 0	Monitored object: CloudTable instance Dimension:	1 minute

Metric ID	Metric	Meaning	Value Range	Monitored Object and Dimension	Monitoring Period (Raw Data)
				instance_id	
rs_writerequestscount	Write Requests	Number of write requests of a RegionServer	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_readrequestscount	Read Requests	Number of read requests of a RegionServer	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_blockcachehitcachingratio	Hit Cache Block Caching Ratio	Block cache hit caching ratio Unit: %	0%~100%	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_blockCacheCountHitPercent	Hit Cache Block Ratio	Block cache hit ratio Unit: %	0%~100%	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_getavgtime	Get Delay (Avg)	Average Get operation delay of the RegionServer per unit time Unit: millisecond	>= 0 ms	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_putavgtime	Put Delay (Avg)	Average Put operation delay of the RegionServer per unit time Unit: millisecond	>= 0 ms	Monitored object: CloudTable instance Dimension: instance_id	1 minute

Metric ID	Metric	Meaning	Value Range	Monitored Object and Dimension	Monitoring Period (Raw Data)
rs_deleteavgtime	Delete Delay (Avg)	Average Delete operation delay of the RegionServer per unit time Unit: millisecond	>= 0 ms	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_getnumops	Get Operations	Number of Get operations of the RegionServer per unit time	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_putnumops	Put Operations	Number of Put operations of the RegionServer per unit time	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_deletenumops	Delete Operations	Number of Delete operations of the RegionServer per unit time	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_queuecalltime_max	RPC Queue Call Time (Max)	Maximum RPC queue call time Unit: millisecond	>= 0 ms	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_queuecalltime_mean	RPC Queue Call Time (Mean)	Mean RPC queue call time Unit: millisecond	>= 0 ms	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_flush_time	Flush	Mean time of flush	>= 0	Monitored	1

Metric ID	Metric	Meaning	Value Range	Monitored Object and Dimension	Monitoring Period (Raw Data)
_mean	Time(Mean)	Unit: millisecond	ms	object: CloudTable instance Dimension: instance_id	minute
rs_compactionqueuesize	Compaction Queue Size	Point in time length of the compaction queue. The number of Stores for compaction in the RegionServer.	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_flushqueueesize	Flush Queue Size	Flush queue size	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_compactioncompletedcount	Compaction Count	Count of compaction	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_flushops_num	Flush Operation Count	Count of flush operation	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_blockcacheevictedcount	Discarded Cache Blocks	Block cache evict count	>= 0	Monitored object: CloudTable instance Dimension: instance_id	1 minute
rs_syncTime_max	Sync WAL Time(Max)	Maximum time it took to sync the WAL to	>= 0 ms	Monitored object:	1 min

Metric ID	Metric	Meaning	Value Range	Monitored Object and Dimension	Monitoring Period (Raw Data)
		HDFS Unit: millisecond		CloudTable instance Dimension: instance_id	ute
rs_syncTime_mean	Sync WAL Time(Mean)	Mean time it took to sync the WAL to HDFS Unit: millisecond	>= 0 ms	Monitored object: CloudTable instance Dimension: instance_id	1 minute

Dimension

Key	Value
instance_id	CloudTable instance

3.9.2 Setting CloudTable Alarm Rules

You can set CloudTable alarm rules to customize the monitored objects and notification policies. Then, you can learn about CloudTable running status in a timely manner.

The CloudTable alarm rules include alarm rule name, instance, metric, threshold, monitoring interval and whether to send notification. This section describes how to set CloudTable alarm rules.

Setting CloudTable Alarm Rules


- Step 1** Log in to the CloudTable management console.
- Step 2** In the left navigation pane, click **Cluster Mode**.
- Step 3** In the upper right corner of the cluster list, enter the name of a cluster in the search box and click .

Figure 3-42 Searching for a cluster by name



Step 4 In the **Operation** column of the cluster, click **Monitor** to access the Cloud Eye management console. Then you can view CloudTable monitoring information.

NOTE

The status of the cluster whose monitoring information you want to view must be **In service**.

Step 5 In the left navigation pane, choose **Alarm Management > Alarm Rules**.

Step 6 On the displayed **Alarm Rule** page, click **Create Alarm Rule** to create an alarm rule. Alternatively, modify an existing alarm rule to set an alarm rule for CloudTable.

This section uses the modification of an existing alarm rule as an example.

In the displayed **Create Alarm Rule** dialog box, set parameters as prompted.

1. Enter an alarm rule name and description.
2. Configure the alarm parameters as prompted.

Figure 3-43 Select the monitoring scope

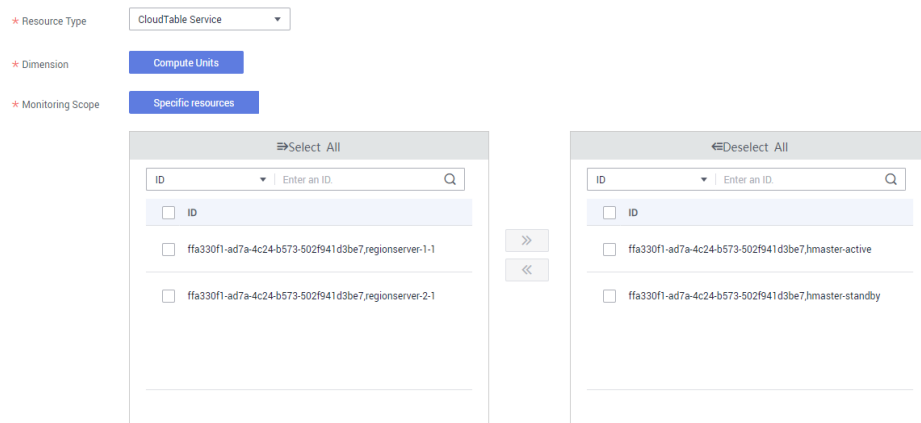


Figure 3-44 Setting the alarm policy





Table 3-19 Alarm rule parameter description

Parameter	Description	Example Value
Resource Type	Name of the cloud service resource for which the alarm rule is configured	CloudTable Service
Dimension	Metric dimension of the alarm rule. Only the Compute Units option is supported.	Computing Units
Monitoring Scope	Resource scope to which an alarm rule applies. Select Specific resources and select one or more monitoring objects. For CloudTable, select the compute node ID of the cluster and click  to synchronize the monitored objects to the dialog box on the right.	Specific resources
Method	Select Use template or Create manually as required. <ul style="list-style-type: none"> If no alarm template is available, set Method to Create manually and configure related parameters to create an alarm rule. If you have available alarm rule templates, set Method to Use template, so that you can use a template to quickly create alarm rules. 	Create manually
Template	This parameter is valid only when Use template is selected. Select the template to be imported. If no alarm template is available, click Create Custom Template to create one that meets your requirements.	-
Alarm Policy	This parameter is valid only when Create manually is selected.	-

Parameter	Description	Example Value
	Set the policy that triggers an alarm. For example, trigger an alarm if the CPU usage equals to or is greater than 80% for 3 consecutive periods. Monitored CloudTable Metrics lists the monitoring metrics supported by CloudTable.	
Alarm Severity	Severity of an alarm. Valid values are Critical , Major , Minor , and Informational .	Major

- Configure the alarm notification parameters as prompted.

Figure 3-45 Configuring alarm notification

Alarm Notification

* Validity Period -

* Notification Object

Create an SMN topic and click refresh to make it available for selection.

* Trigger Condition Generated alarm Cleared alarm

Table 3-20 Alarm notification parameters

Parameter	Description	Example Value
Alarm Notification	Whether to notify users when alarms are triggered. Notifications can be sent as emails or text messages, or HTTP/HTTPS requests sent to the servers. You can enable (recommended) or disable Alarm Notification .	Yes
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If Validity Period is set to 00:00-8:00 , Cloud Eye send notifications only within 00:00-8:00.	-
Notification Object	Name of the topic to which the alarm notification is sent. If you enable Alarm Notification , you need to select a topic. If no desired topics are available, create one first, whereupon the SMN service is invoked. For details about	-

Parameter	Description	Example Value
	how to create a topic, see the <i>Simple Message Notification User Guide</i> .	
Trigger Condition	Condition for triggering the alarm. You can select Generated alarm , Cleared alarm , or both.	-

- After the configuration is complete, click **Next**.
After the alarm rule is created, if the metric data reaches the specified threshold, Cloud Eye will immediately inform you that an exception has occurred.

---End

3.9.3 Viewing Cluster Monitoring Information

Scenario

Cloud Eye on the public cloud monitors CloudTable cluster running status. You can view CloudTable monitoring metrics on the management console. According to the monitoring information, you can quickly learn about cluster health status and key system information.

Monitored data requires some time for transmission and display. Cloud Eye provides CloudTable cluster node monitoring information obtained 5 to 10 minutes before. You can view the monitored data of a newly created CloudTable cluster 5 to 10 minutes later.

Background Information

- Monitoring metrics of an unavailable CloudTable cluster node cannot be displayed on the Cloud Eye page. You can view the monitoring information only after the CloudTable cluster node is restarted or recovered.
- Cloud Eye will delete a CloudTable cluster node that becomes faulty for 1 hour from the monitoring list and will not monitor it any more. However, you need to manually clear its alarm rules.

Procedure



- Step 1** Log in to the CloudTable management console.
- Step 2** Click  in the upper left corner to select a region.
- Step 3** In the left navigation pane, click **Cluster Mode**.
- Step 4** In the upper right corner of the cluster list, enter the name of a cluster in the search box and click .



Figure 3-46 Searching for a cluster



Step 5 In the **Operation** column of the cluster, click **Monitor** to access the Cloud Eye management console. Then you can view CloudTable monitoring information.

Table 3-21 describes parameters about CloudTable service monitoring.


Table 3-21 Instance monitoring parameters

Parameter	Description
ID	Name of the monitored instance
Permanent Data Storage	Name of the OBS bucket where raw data of all monitoring metrics is stored Click  next to ID to expand the instance, and click Configure Data Storage to configure the permanent data storage.
Operation	Click  next to ID to expand the instance. You can see the following operation buttons: View Monitoring Graph , Create Alarm Rule , and Configure Data Storage .

Step 6 In the **Operation** column of the instance, click **View Monitoring Graph** to access the Cloud Eye management console. In the left navigation pane, choose **CloudTable > Monitoring metric**.

For more information about monitoring graphs, see the *Cloud Eye User Guide*.

Step 7 On the **Monitoring metric** page, graphs of raw data collected in **1h**, **3h**, and **12h** are displayed. In the right corner of the monitoring graph of raw metric data, **Max** and **Min** values of the target metric are displayed. By default, nine monitoring graphs are displayed on one page. If the number of metrics of a service instance is more than nine, you can click **Load more** to view more monitoring graphs.

Step 8 In the upper right corner of the monitoring graph, click  to enlarge the graph for viewing detailed data.

Step 9 On the displayed page, select an existing time range or manually specify a time range for instance monitoring.

The system allows you to select a fixed time range or customize the time range.

- The fixed time range can be **1h**, **3h**, **12h**, **1d**, **7d**, or **30d**.
 - If you select **1h**, **3h**, **12h**, or **1d**, raw metric data is displayed by default.
 - If you select **7d** or **30d**, aggregated metric data is displayed by default.
- Custom time ranges can be any time within the last six months.

Step 10 On the page of graph details, click **Settings** to aggregate monitoring data. Select **Enable rollup**. Set **Period** and **Statistic**.

- **Enable rollup** dictates what type of metric data is displayed. If you select **Enable rollup**, aggregated metric data is displayed. If you do not select **Enable rollup**, raw metric data is displayed.
- **Period** indicates a time range during which raw data is collected. The value can be **5 minutes**, **20 minutes**, **1 hour**, **4 hours**, or **1 day**. In the **Period** drop-down, several values are recommended based on your custom time range.
- **Statistic** indicates the rollup method to be used, including **Max.**, **Min.**, **Avg.**, **Sum**, and **Variance**.

----End

3.10 Audit Logs

3.10.1 Supported Key Audit Log Operations

CloudTable uses CTS to record operations associated with CloudTable for later query, audit, and backtrack operations.

The following key operation traces of CloudTable are recorded in audit logs. For details, see [Table 3-22](#).

Table 3-22 CloudTable operation traces supported by CTS

Operation	Trace Name	Resource Type
Creating a cluster	createCluster	cluster
Deleting a cluster	deleteCloudTableCluster	cluster
Expanding cluster capacity	growCloudTableCluster	cluster
Restarting a cluster	rebootCloudTableCluster	cluster
Setting the storage quota	storageClusterAction	cluster
Modifying a feature	modifyClusterFeatures	cluster
Configuring parameters	modifyClusterSetting	cluster

3.10.2 Viewing Audit Logs

After you enable CTS, the system starts recording CloudTable operation logs. Operation records of the last 7 days can be viewed on the CTS management console.

This section includes the following content:

- [Enabling CTS](#)

- [Disabling the Audit Log Function](#)
- [Viewing CTS Logs of CloudTable](#)

Enabling CTS

A tracker will be automatically created after CTS is enabled. All traces recorded by CTS are associated with a tracker. Currently, only one tracker can be created for each account.

Step 1 On the CloudTable management console, choose **Service List > Management & Deployment > Cloud Trace Service**. The CTS management console is displayed.

Step 2 Enabling CTS

If you are a first-time CTS user and do not have any created trackers in the tracker list, enable CTS first. For details, see [Enabling CTS](#) in the *Cloud Trace Service Getting Started*.

If you have enabled CTS, the system has automatically created a management tracker. Only one management tracker can be created and it cannot be deleted. You can also manually create a data tracker. For details, see [Creating a Tracker](#) in the *Cloud Trace Service User Guide*.

---End

Disabling the Audit Log Function

If you want to disable the audit log function, disable the tracker in CTS.

Step 1 On the CloudTable management console, choose **Service List > Management & Deployment > Cloud Trace Service**. The CTS management console is displayed.

Step 2 Disable the audit log function by disabling the tracker. To enable the audit log function again, you only need to enable the tracker.

For details about how to enable or disable a tracker, see [Disabling or Enabling a Tracker](#) in the *Cloud Trace Service Getting Started*.

----End

Viewing CTS Logs of CloudTable

Step 1 Log in to the management console.

Step 2 Click **Service List** on the upper part of the page and choose **Management & Deployment > Cloud Trace Service**.

Step 3 In the navigation pane on the left, choose **Trace List**.

Step 4 In the upper right corner of the trace list, click **Filter** to set the search criteria.

The following four filter criteria are available:

- **Trace Source, Resource Type, and Search By**
 - **Trace Source:** Select **CloudTable**.
 - **Resource Type:** Select **All resource types** or specify a resource type.
 - **Search By:** Select **All** or any of the following options:
 - **Trace name:** If you select this option, you also need to select a specific trace name.

- **Resource ID:** If you select this option, you also need to select or enter a specific resource ID.
 - **Resource name:** If you select this option, you also need to select or enter a specific resource name.
- **Operator:** Select a specific operator (at user level rather than tenant level).
- **Trace Status:** Available options include **All trace statuses**, **normal**, **warning**, and **incident**. You can only select one of them.
- **Start Date and End Date:** You can specify the time period to query traces.

Step 5 Click **Query**.


Step 6 Click  on the left of the trace to be queried to extend its details.

Figure 3-47 Trace

Trace Name	Resource Ty...	Trace So...	Resource ID ⓘ	Resource Name(Trace Stat...	Operator ⓘ	Operation Time	Operation
rebootCloud...	cluster	CloudTab...	689ae304-086...	cloudtable-93...	● normal		05/29/2018 14:50:16 GMT+...	View Trace

Trace ID	8b253930-630c-11e8-82ae-286ed488cbe3	Source IP Address	
Trace Type	ConsoleAction	Generated	05/29/2018 14:50:16 GMT+08:00

Step 7 Locate the row containing the target trace and click **View Trace** in the **Operation** column.

Figure 3-48 Viewing a trace

View Trace ✕

```

{
  "service_type": "CloudTable",
  "user": {
    "name": " ",
    "id": "2f210e62179a425c96e5b57f25d1d492",
    "domain": {
      "name": " ",
      "id": "760dfba8429649e7a477445e1033a839"
    }
  },
  "time": "05/29/2018 14:50:16 GMT+08:00",
  "code": 200,
  "resource_type": "cluster",
  "resource_name": "cloudtable-0000",
  "resource_id": "689ae304-086b-4fa9-984e-16c50e3d75bf",
  "source_ip": " ",
  "trace_name": "rebootCloudTableCluster",
  "trace_type": "ConsoleAction",
  "api_version": "v1.0",
  "record_time": "05/29/2018 14:50:16 GMT+08:00",
  "trace_id": "8b253930- ",
  "trace_status": "normal"
}

```

For details about key fields in the CTS trace structure, see the [Trace Structure](#) in the *Cloud Trace Service User Guide*.

----End

A Change History

Release Date	Description
2019-11-07	<p>This issue is the nineteenth official release.</p> <p>Modified the following section:</p> <ul style="list-style-type: none"> • Monitored CloudTable Metrics
2019-08-16	<p>This issue is the eighteenth official release.</p> <p>Modified the following section:</p> <ul style="list-style-type: none"> • Using HBase Shell to Access a Cluster
2019-05-25	<p>This issue is the seventeenth official release.</p> <ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> – Authorizing Other Users to Access HBase in a Cluster with IAM Authentication Enabled – Creating a User and Granting Permissions • Modified the following sections: <ul style="list-style-type: none"> – Preparing a Local Windows Environment (VPN Connection Mode)
2019-05-08	<p>This issue is the sixteenth official release.</p> <ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> – Preparing a Local Windows Environment (VPN Connection Mode) – OpenTSDB Overview – Connecting to OpenTSDB • Modified the following sections: <ul style="list-style-type: none"> – Preparing an ECS – Using the GeoMesa Command Line Tools
2019-04-04	<p>This issue is the fifteenth official release.</p> <ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> – Introduction to Cluster Management – Expanding Cluster Capacity – Using the HBase Shell for Full-Text Indexing

Release Date	Description
	<ul style="list-style-type: none"> • Modified the following section: <ul style="list-style-type: none"> – Overview of Full-Text Search
2019-03-06	<p>This issue is the fourteenth official release.</p> <ul style="list-style-type: none"> • Added the following section: <ul style="list-style-type: none"> – Overview of Full-Text Search
2019-01-15	<p>This issue is the thirteenth official release.</p> <ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> – Using CopyTable to Import Data – Sample Code for Copying a Table • Modified the following sections: <ul style="list-style-type: none"> – Creating a Cluster Quickly – Creating a Cluster – Viewing Basic Cluster Information – Preparing an ECS
2019-01-04	<p>This issue is the twelfth official release.</p> <ul style="list-style-type: none"> • Modified the following section: <ul style="list-style-type: none"> – Creating a Cluster
2018-12-07	<p>This issue is the eleventh official release.</p> <ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> – Introduction to the Cluster Mode • Modified the following sections: <ul style="list-style-type: none"> – Creating a Cluster – Using HBase Shell to Access a Cluster – Accessing the HBase Web UI • Deleted the following section: <ul style="list-style-type: none"> – Overview
2018-09-25	<p>This issue is the tenth official release.</p> <ul style="list-style-type: none"> • Modified the following sections: <ul style="list-style-type: none"> – Enabling OpenTSDB – Using HBase Shell to Access a Cluster
2018-09-10	<p>This issue is the ninth official release.</p> <ul style="list-style-type: none"> • Modified the following sections: <ul style="list-style-type: none"> – Enabling OpenTSDB – Using HBase Shell to Access a Cluster
2018-08-03	<p>This issue is the eighth official release.</p> <ul style="list-style-type: none"> • Modified the following sections: <ul style="list-style-type: none"> – Creating a Cluster

Release Date	Description
	<ul style="list-style-type: none"> - Using HBase Shell to Access a Cluster
2018-06-30	<p>This issue is the seventh official release.</p> <ul style="list-style-type: none"> • Modified the following sections: <ul style="list-style-type: none"> - Creating a Cluster - Using HBase Shell to Access a Cluster • Deleted the following sections: <ul style="list-style-type: none"> - Application Development Guide: This section is released separately.
2018-05-31	<p>This issue is the sixth official release.</p> <ul style="list-style-type: none"> • Modified the following sections: <ul style="list-style-type: none"> - Using CDM to Migrate Data to CloudTable - Using the Import Tool to Importing Data - Supported Key Audit Log Operations - Viewing Audit Logs
2018-04-26	<p>This issue is the fifth official release.</p> <ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> - Using the GeoMesa Command Line Tools - GeoMesa Command Line • Modified the following sections: <ul style="list-style-type: none"> - Creating a Cluster Quickly
2018-02-12	<p>This issue is the fourth official release.</p> <ul style="list-style-type: none"> • Added the following sections: <ul style="list-style-type: none"> • Monitored CloudTable Metrics • Setting CloudTable Alarm Rules • Viewing Cluster Monitoring Information • Modified the following sections: <ul style="list-style-type: none"> - Creating a Cluster - Preparing an ECS - Using HBase Shell to Access a Cluster
2018-01-03	<p>This issue is the third official release.</p> <p>Added the following sections:</p> <p>Modifying HBase Parameters of the Cluster</p> <p>Modified the following sections:</p> <ul style="list-style-type: none"> • Overview • Creating a Cluster • Viewing Basic Cluster Information
2017-11-29	<p>This issue is the second official release.</p> <p>Modified the following sections:</p>

Release Date	Description
	<ul style="list-style-type: none"><li data-bbox="576 324 820 353">• Creating a Cluster<li data-bbox="576 365 1050 394">• Using HBase Shell to Access a Cluster
2017-09-30	This issue is the first official release.