

Alibaba Cloud

ApsaraDB for HBase Operation and Maintenance Guide

Document Version: 20201026

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions









Style	Description	Example
 Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	 Danger: Resetting will result in the loss of user configuration data.
 Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	 Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
 Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: If the weight is set to 0, the server no longer receives new requests.
 Note	A note indicates supplemental instructions, best practices, tips, and other content.	 Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type .
Bold	Bold formatting is used for buttons, menus, page names, and other UI elements.	Click OK .
Courier font	Courier font is used for commands	Run the <code>cd /d C:/window</code> command to enter the Windows system folder.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <i>Instance_ID</i>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

Table of Contents

1. Access a web UI	05
2. Scale out a cluster	07
3. Delete protection	08
4. Configure a whitelist or a security group	09
5. Connectivity test	11
6. Use RAM users to manage ApsaraDB for HBase clusters	13
7. HBase Shell	14
8. Tag management	17
8.1. Create a tag	17
8.2. Unbind a tag	17
8.3. Use tags to filter ApsaraDB for HBase clusters	18

1. Access a web UI

This topic describes how to access open source components after you create an ApsaraDB for HBase cluster. You can access components such as HBase, Ganglia, HDFS, and ClusterManager (if you use ApsaraDB for HBase Performance-enhanced Edition).

1. Log on to the [ApsaraDB for HBase console](#), find your cluster, and click **Manage**.
2. In the left-side navigation pane, choose **Database Connection**. You can find the web UIs of the components in the **UI Access** section.
3. If this is the first time that you are accessing the open source components, you must set the username and password. If you have already set the username and password, skip this step.
 - i. Click **Reset UI Access Password** on the right side.
 - ii. Enter a username and password in the dialog box that appears. The username and password must be 2 to 30 characters in length and can contain letters, digits, hyphens (-), and underscores (_).
 - iii. Click **OK** to set the username and password.
 - iv. You can repeat the preceding steps to change the username and password. Only one username and password pair is supported. Therefore, the new username and password will overwrite the current ones.
4. Configure the IP whitelist. For more information, see [Configure the IP whitelist](#).
5. After you complete the preceding steps, click the hyperlink of the open source component, and enter the username and password to access the web UI.

Note: Connections to web UIs are established over HTTPS. The system automatically generates a unique certificate for each cluster. The certificate is not authorized. If the "Your connection is not private" message is prompted, click **Advanced** and then click **Proceed** to access the web UI.

Your connection is not private

Attackers might be trying to steal your information from **hb-bp190xuaw7868bec5proxy-003.hbase.rds.aliyuncs.com** (for example, passwords, messages, or credit cards). [Learn more](#)

NET::ERR_CERT_AUTHORITY_INVALID

☐ Automatically send some [system information and page content](#) to Google to help detect dangerous apps and sites. [Privacy policy](#)

HIDE ADVANCED

Back to safety

This server could not prove that it is **hb-bp190xuaw7868bec5proxy-003.hbase.rds.aliyuncs.com**; its security certificate is not trusted by your computer's operating system. This may be caused by a misconfiguration or an attacker intercepting your connection.

[Proceed to hb-bp190xuaw7868bec5proxy-003.hbase.rds.aliyuncs.com \(unsafe\)](#)

HBase UI:

⚠ Not Secure

https://hb-bp190xuaw7868bec5proxy-003.hbase.rds.aliyuncs.com/master-1/master-status?spm=5176.8772746.0.0.311227c8UHambw

HBase

Home

Logs

Log Level

Debug Dump

Metrics Dump

Conf

Jstack

Connections

TopRegions

Replication Topology

Attribute Name	Value	Description
AliHB Version	1.4.3.4, revision=f384f7769a602033739a3760e8044189180827e5	Alihb is a HBase branch originated from Alibaba Group where the versions of alihb-1.x are based on official apache hbase-1.1
AliHB Compiled	Tue Aug 8 20:43:56 CST 2017, admin	When AliHB version was compiled and by whom
AliHB Source Checksum	886f4202f9e4fea2af611f1642f84a08	AliHB source MD5 checksum
Hadoop Version	2.7-adh6u2.2.3, revision=381541	Hadoop version and revision
Hadoop Compiled	2017-01-04T10:16Z, wangxiang	When Hadoop version was compiled and by whom

2.Scale out a cluster

In the ApsaraDB for HBase console, you can scale out the nodes of a cluster or expand the storage space of the cluster

□

To add a node to a cluster, you must ensure that the new node and existing nodes are the same. They must be of the same specification and use the same type of disks. An ApsaraDB for HBase cluster can contain up to 100 nodes. You can add up to 5 nodes to a cluster at a time. If you want to increase the limit, submit a ticket.

When you add more nodes to a cluster, you also add more disks to the cluster. This is equivalent to expanding your storage space.

Adding a node to a cluster does not pose any negative impacts on the cluster.

Expand storage space

You can directly expand the storage space of the disks used by a cluster. 50 GB is the minimum size for each time you expand the storage space. Each disk can be expanded to up to 2,000 GB.

After you expand the storage space, a rolling restart is launched for the RegionServers and DataNodes. This only has minor impacts on the cluster.


3.Delete protection

To prevent important clusters from being deleted by mistake, you can enable delete protection in the ApsaraDB for HBase console.

Procedure

1. Log on to the [ApsaraDB for HBase console](#).
2. Select the region where the cluster is deployed.
3. Find the target cluster and click the cluster ID.
4. In the left-side navigation pane, click **Basic Information**.
5. Click **Enable of Delete Protection**.

6. In the dialog box that appears, click **OK**.

 **Note** After you enable delete protection for a cluster, you cannot delete the cluster in the console or by calling API operations. To delete a cluster, you must disable delete protection.

4. Configure a whitelist or a security group

After you create an ApsaraDB for HBase cluster, you must configure a whitelist or add Elastic Compute Service (ECS) security groups for the cluster. This allows the clients included in a whitelist or security group to access the cluster.

Prerequisites

When you add an ECS security group as a whitelist for an ApsaraDB for HBase cluster, the ECS instance and the ApsaraDB for HBase cluster must be deployed in the same Virtual Private Cloud (VPC) network.

Context

To ensure data security, access to a newly created ApsaraDB for HBase cluster is not allowed by default:

- You are not allowed to access open source components in the cluster, such as the HBase, Ganglia, and HDFS components.
- You are not allowed to read or write data in the ApsaraDB for HBase cluster.

Before you connect to the cluster, you must add the IP addresses of your clients to the whitelist.

Configure a whitelist

1. Log on to the [ApsaraDB for HBase console](#).
2. Select the region where the cluster is deployed.
3. Find the target cluster and click the cluster ID.
4. In the left-side navigation pane, click **Access Control**.
5. Click the **Whitelist Setting** tab and click **Edit Whitelist**.

6. In the Edit Whitelist dialog box, enter the IP addresses or CIDR blocks that are used to access the cluster and click **OK**.

Note

- By default, the whitelist contains only the IP address 127.0.0.1. This indicates that no client is allowed to access the ApsaraDB for HBase cluster.
- If you set Whitelist to 0.0.0.0/0 or leave it blank, all IP addresses are allowed to access your ApsaraDB for HBase cluster. To ensure data security, we recommend that you do not use the 0.0.0.0 IP address or 0.0.0.0/0 CIDR block.

- If you want to use a public IP address to access open source components, enter the public IP address.
- You can also enter a CIDR block. For example, you can enter 192.168.0.0/24 to specify all IP addresses in the subnet 192.168.0.X.

- If you enter more than one IP address or CIDR block, you must separate them with commas (,). Do not add spaces before or after the commas. For example:
192.168.0.1,172.16.213.0/24.

Add security groups

A security group is a virtual firewall that is used to control inbound and outbound traffic of ECS instances in the security group. After a security group is added for an ApsaraDB for HBase cluster, the ECS instances in the security group can access the ApsaraDB for HBase cluster.

- ApsaraDB for HBase Standard Edition and ApsaraDB for HBase Enhanced Edition are supported.
- You can configure both the IP address whitelist and security groups for a cluster. All IP addresses in the IP address whitelist and all ECS instances in security groups are allowed to access the ApsaraDB for HBase cluster.
- You can add up to three security groups for a cluster.

1. Log on to the [ApsaraDB for HBase console](#).
2. Select the region where the cluster is deployed.
3. Find the target cluster and click the cluster ID.
4. In the left-side navigation pane, click **Access Control**.
5. Click the **Security Group** tab and click **Add Security Group**.

6. In the Add Security Group dialog box, select the security groups and click OK.

5.Connectivity test

For security purposes, ApsaraDB for HBase does not support the ICMP protocol.

The ping command sends ICMP packets to test the connectivity. Therefore, use the telnet command instead of the ping command to test the connectivity to your ApsaraDB for HBase cluster.

ApsaraDB for HBase allows you to connect to your cluster over a public or internal network. To connect to your cluster over a public network, you must first apply for a public endpoint for your cluster. **Note:**We recommend that you connect to your ApsaraDB for HBase cluster over a public network only for development purposes. Alibaba Cloud does not guarantee the response speed of production environments over a public network.

- Select one of the ZooKeeper addresses provided in the ApsaraDB for HBase console.
- Run the telnet command to test the connectivity to port 2181. You can also change the address and port number to test the connectivity to Thrift and other database engines.
 - The following prompt is displayed if you are connected to ApsaraDB for HBase:

```
[xx@yy-MacBook-Pro ~]$ telnet hb-xxxxx-001.hbase.rds.aliyuncs.com 2181
Trying 10.10.10.10...
Connected to hb-xxxxx-001.hbase.rds.aliyuncs.com .
Escape character is '^]'.
```

- The following prompt is displayed if you fail to connect to ApsaraDB for HBase:

```
[[xx@yy-MacBook-Pro ~]$ telnet hb-xxxxx-001.hbase.rds.aliyuncs.com 2181
Trying 10.10.10.10...
telnet: connect to address 10.10.10.10: Operation timed out
telnet: Unable to connect to remote host
```

Troubleshooting

Possible causes:

- You have not added your client to the whitelist of ApsaraDB for HBase. You must add your client to the whitelist before you can connect to ApsaraDB for HBase over a public or internal network. You can log on to the ApsaraDB for HBase console and configure the `whitelist`. For more information, see [Configure the whitelist](#).
- ApsaraDB for HBase is not allowed to access your local network. For example, if you use an Elastic Compute Service (ECS) instance, you must allow ApsaraDB for HBase to access the IP address and port of your ECS instance.
- You cannot connect to a VPC network from a classic network.
- Your instance and ApsaraDB for HBase cluster are connected to two VSwitches in a VPC network. In most cases, this means that you are trying to connect to the cluster across zones. To do this, you must add routes.
- You are connecting to ApsaraDB for HBase across VPC networks or regions. By default, VPC networks are isolated from each other. If you want to connect to ApsaraDB for HBase across

VPC networks, use [Express Connect](#) .

- You are connecting to ApsaraDB for HBase from a network outside Alibaba Cloud, such as the private network of your enterprise.
 - Solution 1: apply for a public endpoint for your ApsaraDB for HBase cluster. For more information, see [Connect to ApsaraDB for HBase from a public network](#)
 - Solution 2: establish a [leased line](#) .
- You cannot connect to the public offering of Alibaba Cloud from the finance cloud.
- You use an ECS instance to connect to the public endpoint of your ApsaraDB for HBase cluster, but the ECS instance does not have access to the public network. In this case, connect to the VPC endpoint of your cluster.

If the problem remains unsolved after you check for the preceding causes, consult in the [ApsaraDB for HBase Q&A DingTalk group](#).

6. Use RAM users to manage ApsaraDB for HBase clusters

ApsaraDB for HBase allows you to use RAM users to manage clusters. We recommend that you use authorized RAM users to manage clusters to ensure data security.

Procedure

1. Log on to the **RAM console**, and switch to the old console version. In the left-side navigation pane, choose **Users** and create a RAM user.
2. Click the name of the RAM user that you create. In the left-side navigation pane, choose **User Authorization Policies**. Click **Edit Authorization Policy**, and enter keyword **hbase** into the search box in the dialog box that appears.

□

3. Select a permission policy. The **AliyunHBaseReadOnlyAccess** permission policy only grants RAM users the read permission on clusters. The RAM users can view cluster information but cannot perform other operations such as scaling or restart. The **AliyunHBaseFullAccess** permission policy grants RAM users the full management permission on ApsaraDB for HBase clusters.

4. To allow the RAM user to view monitoring data of ApsaraDB for HBase, you must authorize the RAM user to access CloudMonitor. Enter keyword **cloudmonitor** into the search box, select a CloudMonitor permission policy as needed. Two CloudMonitor permission policies are available. They grant RAM users read-only and full management permissions separately.

□

5. After you select permission policies, click **OK**. The RAM user then has the permissions to access ApsaraDB for HBase and CloudMonitor. You can use the RAM user to log on to the ApsaraDB for HBase console and manage clusters.

7.HBase Shell

This topic describes basic HBase Shell commands.

If you are using ApsaraDB for HBase Enterprise Edition, see [Configure the HBase Shell](#) to configure the basic environment.

If you are using ApsaraDB for HBase Performance-enhanced Edition, see [Access the HBase Shell](#) to configure the basic environment.

1. Connect to your ApsaraDB for HBase cluster.

Navigate to the bin folder of HBase and run the following command to start the HBase Shell.

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

After you start the HBase Shell, the following prompt is displayed. You can then run HBase Shell commands.

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

2. Display HBase Shell help.

Run the following command to list basic HBase Shell commands and information about how to use these commands.

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

3. Create a table.

Run the create command to create a table. You must specify the table name and ColumnFamily name when you create a table.

```
hbase(main):001:0> create 'test', 'cf'
0 row(s) in 0.4170 seconds

=> Hbase::Table - test
```

4. Query information about a table.

Run the list command to query information about a table.

```
hbase(main):002:0> list 'test'
TABLE
test
1 row(s) in 0.0180 seconds

=> ["test"]
```

5. Insert data into a table.

Run the **put** command to insert a row into an HBase table.

```
hbase(main):003:0> put 'test', 'row1', 'cf:a', 'value1'
0 row(s) in 0.0850 seconds

hbase(main):004:0> put 'test', 'row2', 'cf:b', 'value2'
0 row(s) in 0.0110 seconds

hbase(main):005:0> put 'test', 'row3', 'cf:c', 'value3'
0 row(s) in 0.0100 seconds
```

In this example, three rows are inserted into a table. **rowx** specifies the primary key (rowkey) of a row to be inserted. **cf:x** specifies a custom column. The number of custom columns are not limited. In this example, three columns are specified. Values a, b, and c are qualifiers, which refer to column names.

6. Query all data in a table.

The **scan** command can be used to query HBase data. You can use this command to scan a table or query data within the specified range. However, this command returns the query results slower compared with the **get** command. In this example, this command is used because the demo database stores only a small amount of data.

```
hbase(main):006:0> scan 'test'

ROW                COLUMN+CELL
row1                column=cf:a, timestamp=1421762485768, value=value1
row2                column=cf:b, timestamp=1421762491785, value=value2
row3                column=cf:c, timestamp=1421762496210, value=value3
3 row(s) in 0.0230 seconds
```

7. Query a single row.

Run the **get** command to query a single row.

```
hbase(main):007:0> get 'test', 'row1'

COLUMN            CELL
cf:a               timestamp=1421762485768, value=value1
1 row(s) in 0.0350 seconds
```

8. Disable a table.

You must disable a table before you can delete it or change its settings. Run the **disable** command to disable a table and run the **enable** command to enable a table.

```
hbase(main):008:0> disable 'test'  
0 row(s) in 1.1820 seconds
```

```
hbase(main):009:0> enable 'test'  
0 row(s) in 0.1770 seconds
```

9. Delete a table.

Run the **drop** command to delete a table. Exercise caution when you use this command.

```
hbase(main):011:0> drop 'test'  
0 row(s) in 0.1370 seconds
```

10. Exit the HBase Shell.

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

8.Tag management

8.1. Create a tag

If you have created a large number of clusters, you can classify these clusters by binding tags to them. Each tag consists of a key and value. You can use a combination of keys and values to classify your clusters into subcategories.

Limits

- You can bind no more than 20 tags to each cluster. Tag keys must be unique. Two or more tags with the same key will overwrite each other.
- You can bind tags to up to 50 clusters at a time.
- Clusters deployed in different regions do not share the same tag namespace.
- After you unbind a tag, if this tag is not bound to any other clusters, the tag is deleted.

Procedure

1. Log on to the [ApsaraDB for HBase console](#).
2. In the upper-left corner of the page, select the region where your ApsaraDB for HBase clusters are deployed.
 -
3. Select a mode to create tags.
 - Create a tag: After you select a cluster, select **More > Edit Tag** in the Actions column.
 - Create one or more tags: Select the clusters that you want to create tags for, and click **Edit Tag**, as shown in the following figure.
 -
4. Click **Create**, enter a **Key** and **Value**, and click **OK**, as shown in the following figure.

Note: If you have already created some tags, you can select **Existing Tags**.

□

5. Click **OK** to bind the tags to the clusters.

8.2. Unbind a tag

If your cluster no longer needs a tag, you can unbind this tag from the cluster.

Limits

- You can unbind no more than 20 tags at a time.
- After you unbind a tag, if this tag is not bound to any other clusters, the tag is deleted.

Procedure

1. Log on to the [ApsaraDB for HBase console](#).
2. In the upper-left corner of the page, select the region where your ApsaraDB for HBase cluster is deployed.



3. Find your cluster and select **Edit Tag** in the Actions column.
4. Click the X icon next to a tag to delete the tag, as shown in the following figure.



5. Click OK.

8.3. Use tags to filter ApsaraDB for HBase clusters

After you bind tags to ApsaraDB for HBase clusters, you can use these tags to filter ApsaraDB for HBase clusters.


Procedure

1. Log on to the [ApsaraDB for HBase console](#).
2. In the upper-left corner of the page, select the region where your ApsaraDB for HBase clusters are deployed.



3. Click **Tags**, and select a pair of **Key** and **Value** to filter clusters, as shown in the following figure.



 **Note** To delete a filter condition, you can click the X icon next to a tag.