ALIBABA CLOUD

# Alibaba Cloud

## ApsaraDB for Cassandra

## User guide

Document Version: 20210630

Alibaba Cloud

# 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:<br>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:<br>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:<br>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:<br>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 `cd /d C:/window` command to enter the Windows system folder. |
| *Italic* | Italic formatting is used for parameters and variables. | `bae log list --instanceid`<br>*Instance_ID* |
| [] or [a\|b] | This format is used for an optional value, where only one item can be selected. | `ipconfig [-all\|-t]` |
| {} or {a\|b} | This format is used for a required value, where only one item can be selected. | `switch {active\|stand}` |

# Table of Contents

# 1.Change a password

If you forget your password, need to change your password, or have not set a password for an instance, you can set a new password for the instance.

## Instructions

- The default superuser username is cassandra. You can use this account to create other accounts and grant data management permissions as needed.
- After changing the password, you must wait at least 30 seconds before you can log on with the new password, because the old password is cached for 30 seconds before it expires.

1. Log on to the ApsaraDB for Cassandra console. Click a cluster ID to go to the Cluster Details page of the cluster. In the left-side navigation pane, click **Accounts**.

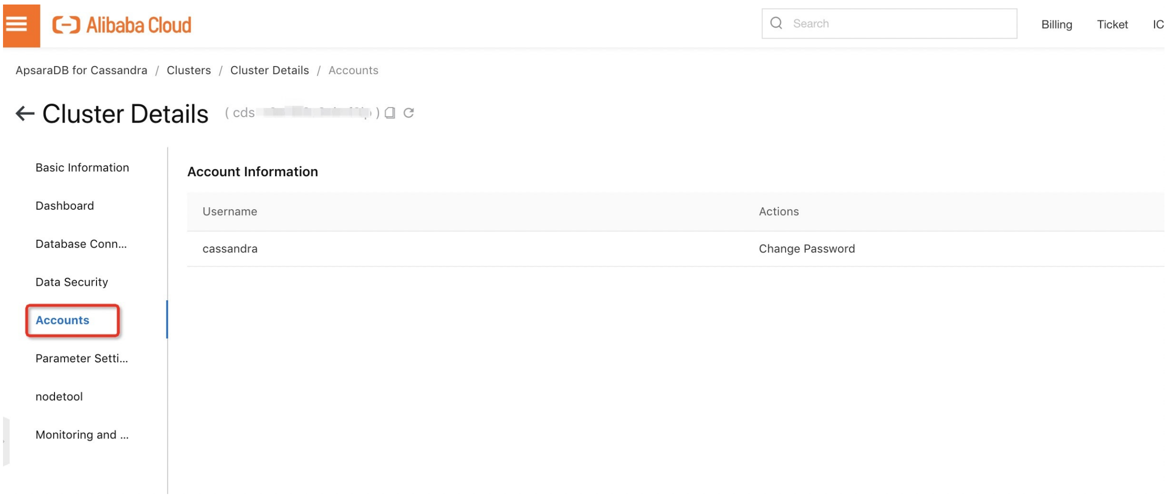

ApsaraDB for Cassandra / Clusters / Cluster Details / Accounts

← Cluster Details ( cds [redacted] )

Account Information

| Username | Actions |
|---|---|
| cassandra | Change Password |

2. Click **Change Password**.

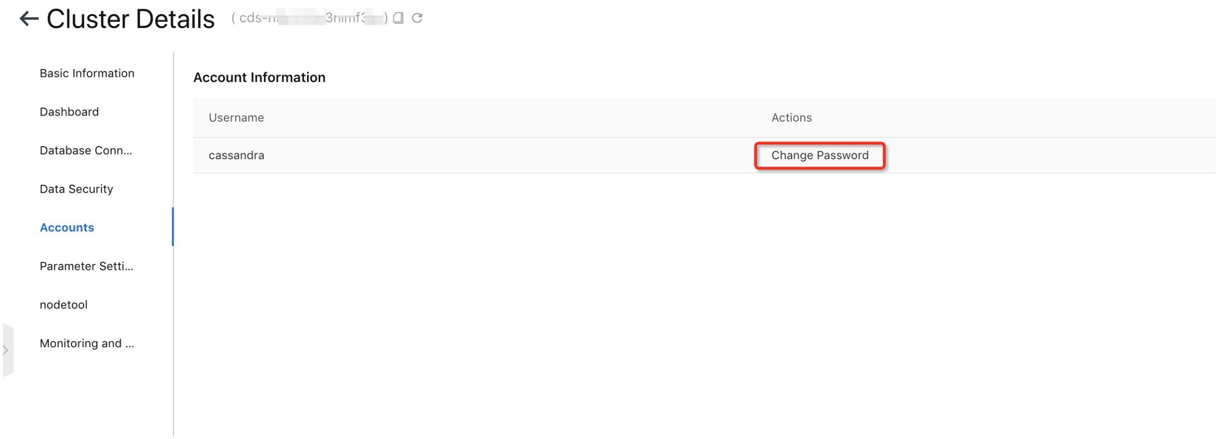
← Cluster Details ( cds-n [redacted] 3nmf3 [redacted] )

Account Information

| Username | Actions |
|---|---|
| cassandra | Change Password |

3. Enter a password as prompted.

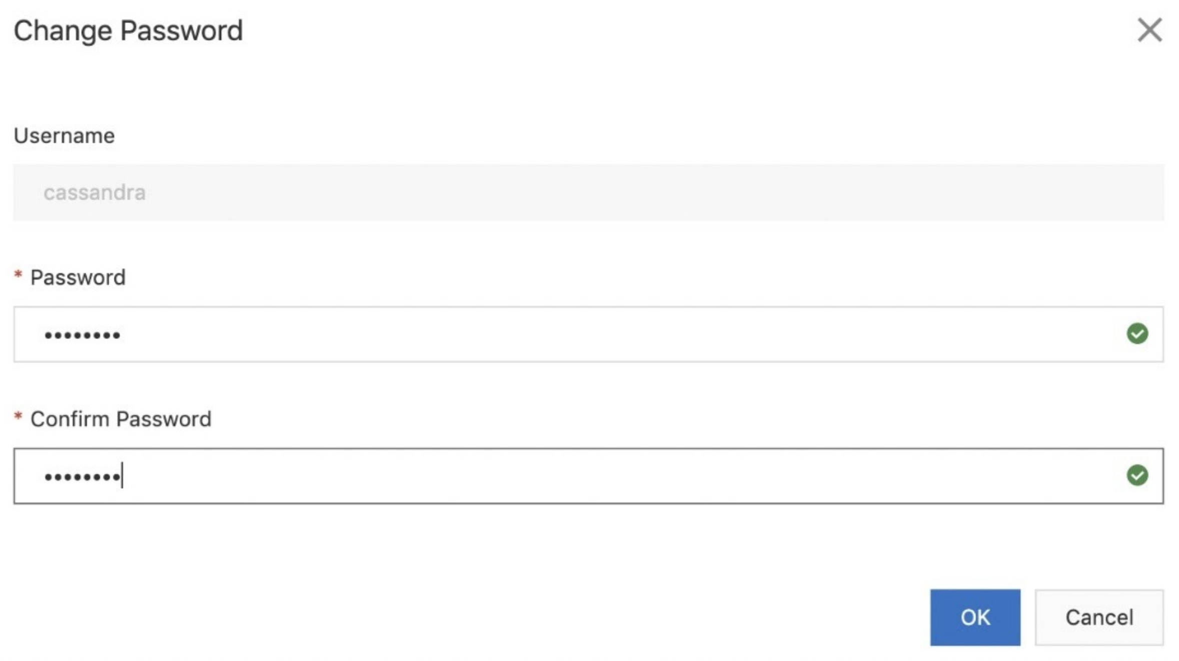

Change Password

Username

cassandra

* Password

••••••••

* Confirm Password

••••••••

# 2.Data migration

## 2.1. Migrate data by running Copy commands

This topic describes how to run Copy commands to migrate data. For more information about Copy commands, see Copy commands.

### Procedure

1. To export data from the source cluster to a CSV file, connect to the source cluster and execute the following statement in cqlsh:

```
COPY table_name [( column_list )]
TO 'file_name'[, 'file2_name', ...] ;
```

After the statement is executed, the data is exported to the file that is specified in the statement.

2. To import the CSV file to the target cluster, connect to the target cluster and execute the following statement in cqlsh:

```
COPY table_name [( column_list )]
FROM 'file_name'[, 'file2_name', ...] ;
```

After the statement is executed, the data is imported to the target cluster.

# 3.Use BulkLoad to import data

## Before you begin

This tool uses a file streaming interface to import data to an ApsaraDB for Cassandra cluster. BulkLoad is one of the fastest ways to migrate offline data to a Cassandra cluster. Before you import data, make the following preparations:

- Create a Cassandra cluster.
- Prepare offline data in SSTable or CSV format.
- Create an independent ECS instance in the same VPC as the Cassandra cluster, and configure security group rules to ensure that the ECS instance can access the Cassandra cluster.

## 1. Create an ECS instance of the client in the same VPC as the Cassandra cluster

We recommend that you create an ECS instance independent of the Cassandra cluster. Otherwise, online services may be affected.

## 2. Create a schema

```
$ cqlsh -f schema.cql  -u USERNAME -p PASSWORD [host]
```

## 3. Prepare data

## 3.1 SSTable data format

Organize a directory in the data/${keyspace}/${table} format and store SSTable data in the directory, as shown in the following example:

```
ls /tmp/quote/historical_prices/
md-1-big-CompressionInfo.db md-1-big-Data.db     md-1-big-Digest.crc32    md-1-big-Filter.db    md-1-big-
Index.db     md-1-big-Statistics.db    md-1-big-Summary.db    md-1-big-TOC.txt
```

In the preceding example, the keyspace parameter is set to quote and the table parameter is set to historical_prices.

## Import data

Run the sstableloader command to specify the data catalog data/${ks}/${table} in the bin directory of the Cassandra distribution.

```
${cassandra_home}/bin/sstableloader -d <ip address of the node> data/${ks}/${table}
```

After the SSTable data is imported, run the following command to check the data: bin/cqlsh -u USERNAME -p PASSWORD [host]

```
$ bin/cqlsh
cqlsh> select * from quote.historical_prices;

 ticker | date                            | adj_close | close     | high      | low       | open      | volume
--------+---------------------------------+-----------+-----------+-----------+-----------+-----------+--------
   ORCL | 2019-10-29 16:00:00.000000+0000 | 26.160000 | 26.160000 | 26.809999 | 25.629999 | 26.600000 | 181000
   ORCL | 2019-10-28 16:00:00.000000+0000 | 26.559999 | 26.559999 | 26.700001 | 22.600000 | 22.900000 | 555000
```

## 3.2 CSV data format

You must first convert CSV data to the SSTable format. Cassandra provides the CQLSSTableWriter tool for generating SSTables. This tool allows you to convert data in a format into the SSTable format. CSV data must also be organized in advance. Therefore, you must compile and run the code for parsing CSV data on your own. The following sample code demonstrates how to use this tool. For more information about this tool, visit the GitHub repository.

```
    // Prepare SSTable writer
    CQLSSTableWriter.Builder builder = CQLSSTableWriter.builder();
    // set output directory
    builder.inDirectory(outputDir)
        // set target schema
        .forTable(SCHEMA)
        // set CQL statement to put data
        .using(INSERT_STMT)
        // set partitioner if needed
        // default is Murmur3Partitioner so set if you use different one.
        .withPartitioner(new Murmur3Partitioner());
    CQLSSTableWriter writer = builder.build();
    /TODO: Read a CSV file. Read each line of a CSV file in an iterative manner.
    while ((line = csvReader.read()) ! = null)
        {
          writer.addRow(ticker,
                  DATE_FORMAT.parse(line.get(0)),
                  new BigDecimal(line.get(1)),
                  new BigDecimal(line.get(2)),
                  new BigDecimal(line.get(3)),
                  new BigDecimal(line.get(4)),
                  Long.parseLong(line.get(6)),
                  new BigDecimal(line.get(5)));
        }
        writer.close();
```

After you generate SSTable data by using the custom program, import the data as described in section 3.1.

# 4.Configure monitoring and alerting functions

ApsaraDB for Cassandra provides monitoring and alerting functions based on CloudMonitor. For more information about the common functions of CloudMonitor, see CloudMonitor documentation.

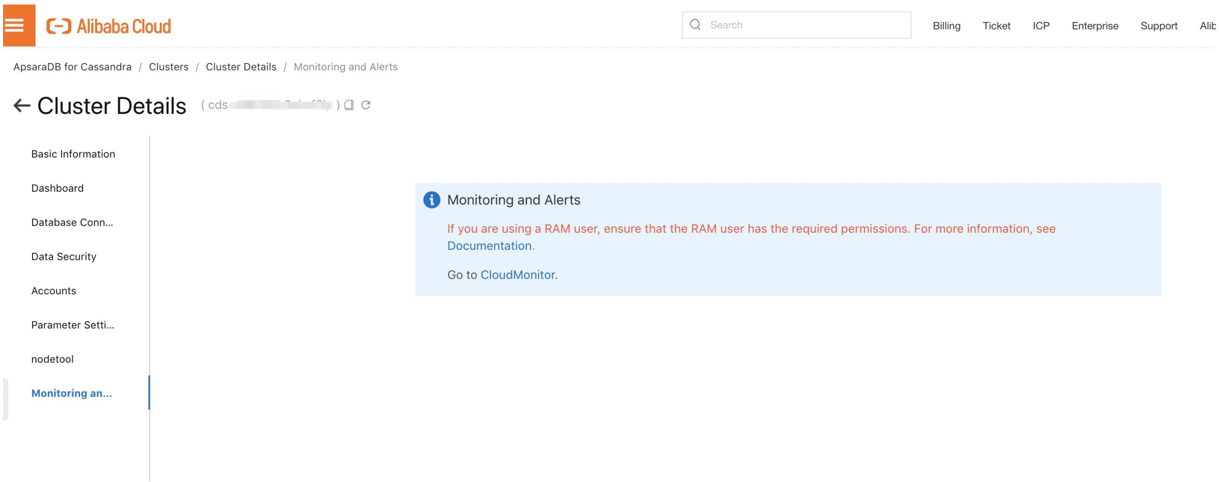

ApsaraDB for Cassandra / Clusters / Cluster Details / Monitoring and Alerts

← Cluster Details ( cds[redacted] )

Monitoring and Alerts

If you are using a RAM user, ensure that the RAM user has the required permissions. For more information, see Documentation.

Go to CloudMonitor.

## Authorize ApsaraDB for Cassandra to access cloud resources

If you access CloudMonitor from ApsaraDB for Cassandra for the first time, you must authorize ApsaraDB for Cassandra to access your CloudMonitor resources. ApsaraDB for Cassandra needs to create a Cassandra group in CloudMonitor and add the current cluster instance to the group.

**Note that you typically need to use an Alibaba Cloud account to perform this authorization.**

## Authorize RAM users to access cloud resources

If you are using a RAM user, ensure that the RAM user is granted the following permission policy.

## CloudMonitor permissions

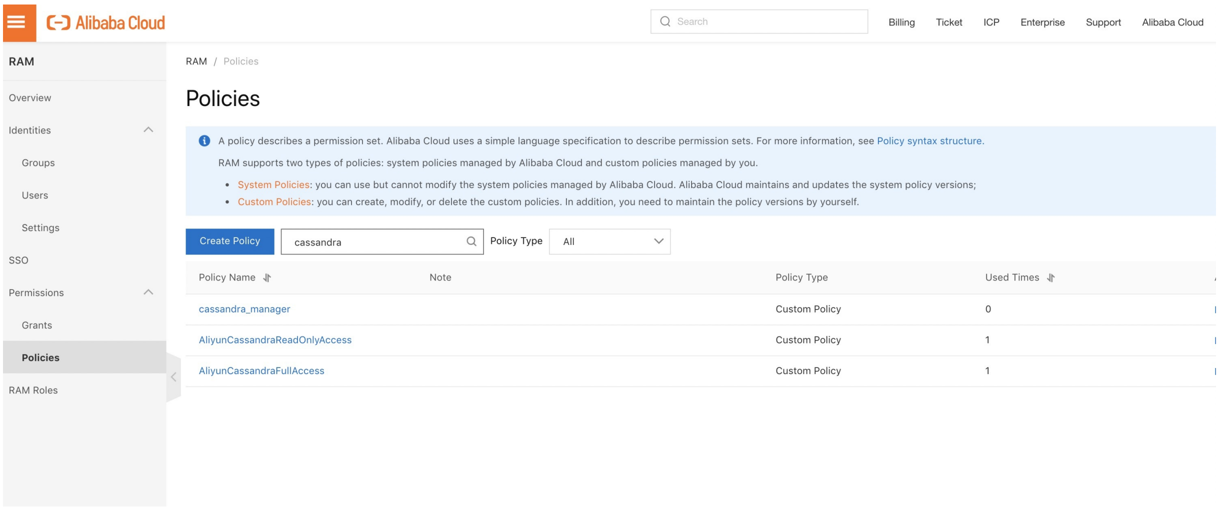

RAM / Policies

Policies

A policy describes a permission set. Alibaba Cloud uses a simple language specification to describe permission sets. For more information, see Policy syntax structure.

RAM supports two types of policies: system policies managed by Alibaba Cloud and custom policies managed by you.

- System Policies: you can use but cannot modify the system policies managed by Alibaba Cloud. Alibaba Cloud maintains and updates the system policy versions;
- Custom Policies: you can create, modify, or delete the custom policies. In addition, you need to maintain the policy versions by yourself.

Create Policy | cassandra | Policy Type: All

| Policy Name | Note | Policy Type | Used Times |
|---|---|---|---|
| cassandra_manager | | Custom Policy | 0 |
| AliyunCassandraReadOnlyAccess | | Custom Policy | 1 |
| AliyunCassandraFullAccess | | Custom Policy | 1 |

## CloudMonitor dashboard

If the relevant permissions are granted, you can go to the CloudMonitor dashboard page.

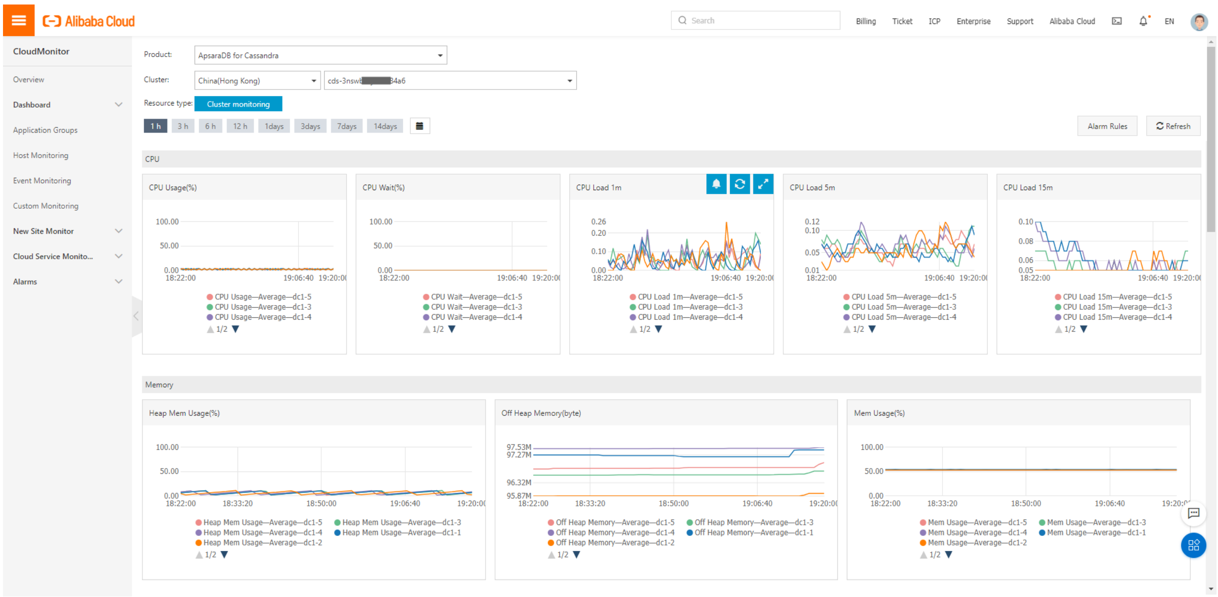

Product: ApsaraDB for Cassandra

Cluster: China(Hong Kong) cds-3nsw[redacted]84a6

Resource type: Cluster monitoring

# 5.Manage RAM users

ApsaraDB for Cassandra also provides common access control features. For more information about how to create RAM users and grant permissions to the users, see Resource Access Management documentation. The following section describes the permission features that are related to ApsaraDB for Cassandra.

## System permission policies

As with most other instance-type cloud services, ApsaraDB for Cassandra provides two default system permission policies:

- AliyunCassandraReadOnlyAccess: grants read-only permissions on an ApsaraDB for Cassandra instance. RAM users that have such permissions cannot perform O&M operations such as scaling up disks, scaling out nodes, enabling and disabling Internet access, and deleting clusters.
- AliyunCassandraFullAccess: grants both read and write permissions on an ApsaraDB for Cassandra instance. This indicates that all the permissions are granted.

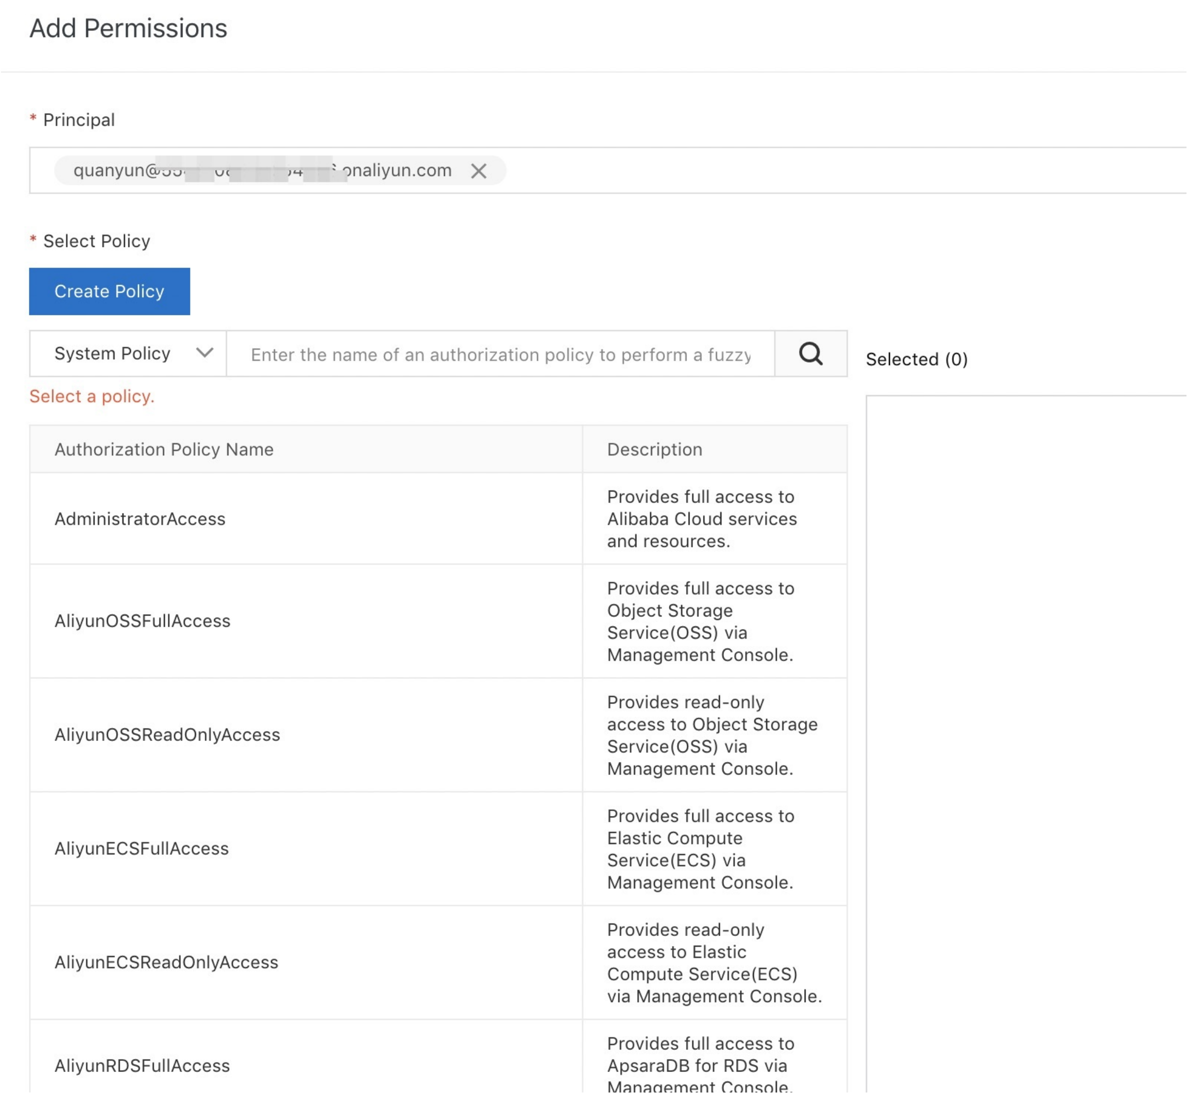

Add Permissions

* Principal

quanyun@[redacted]onaliyun.com ✕

* Select Policy

Create Policy
System Policy
Enter the name of an authorization policy to perform a fuzzy
Selected (0)

Select a policy.

| Authorization Policy Name | Description |
|---|---|
| AdministratorAccess | Provides full access to Alibaba Cloud services and resources. |
| AliyunOSSFullAccess | Provides full access to Object Storage Service(OSS) via Management Console. |
| AliyunOSSReadOnlyAccess | Provides read-only access to Object Storage Service(OSS) via Management Console. |
| AliyunECSFullAccess | Provides full access to Elastic Compute Service(ECS) via Management Console. |
| AliyunECSReadOnlyAccess | Provides read-only access to Elastic Compute Service(ECS) via Management Console. |
| AliyunRDSFullAccess | Provides full access to ApsaraDB for RDS via Management Console. |

# 6.Configure a whitelist

The default whitelist for an ApsaraDB for Cassandra instance only contains the default IP address 127.0.0.1, which means that no device can access the instance. This can ensure the security and stability of the instance. Before using an ApsaraDB for Cassandra instance, you must add IP addresses or Classless Inter-Domain Routing (CIDR) blocks that you use to access databases to the instance whitelist. A properly configured whitelist can guarantee the highest-level security protection for your ApsaraDB for Cassandra instance. We recommend that you maintain the whitelist on a regular basis.

1. Log on to the ApsaraDB for Cassandra console.
2. Click a cluster ID to go to the Cluster Details page of the cluster. In the left-side navigation pane, click **Data Security**. Then, click the **Whitelist Settings** tab.
3. Click Edit Whitelist. In the dialog box that appears, enter the IP addresses that you want to include in the whitelist.

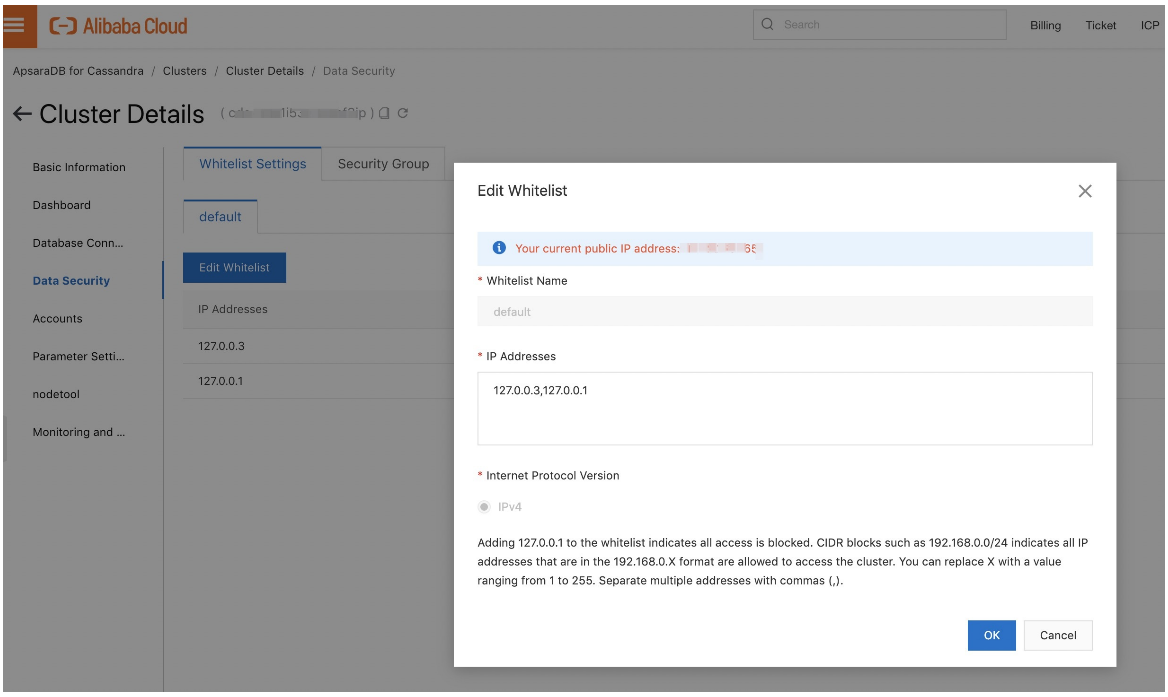

Cluster Details ( [redacted] )

| IP Addresses |
|---|
| 127.0.0.3 |
| 127.0.0.1 |

Edit Whitelist

Your current public IP address: [redacted]

* Whitelist Name: default

* IP Addresses: 127.0.0.3,127.0.0.1

* Internet Protocol Version: ● IPv4

Adding 127.0.0.1 to the whitelist indicates all access is blocked. CIDR blocks such as 192.168.0.0/24 indicates all IP addresses that are in the 192.168.0.X format are allowed to access the cluster. You can replace X with a value ranging from 1 to 255. Separate multiple addresses with commas (,).

- Separate multiple IP addresses with commas (,). Specify IP addresses in the 0.0.0.0/0, 10.23.12.24, or 10.23.12.24/24 format. For CIDR formats such as 10.23.12.24/24, /24 indicates the length of network prefix in the address. The length of network prefix ranges from 1 to 32.
- If you need to access the instance from the Internet, the dialog box shows the public IP address of your device.

# 7.One-touch upgrade

ApsaraDB for Cassandra supports one-touch upgrade in the console. This enables you to upgrade your database version in a short time.

## Procedure

> **Notice** In the version upgrade, each cluster node is serially upgraded. For some version upgrades, cluster restart is required before taking the upgrade into effect, which affects the use of the clusters.
>
> - High availability configuration cluster: If a single node is serially upgraded, the cluster works as expected. This indicates that the high availability configuration cluster supports online version upgrade and the cluster service is not interrupted.
> - Non-high availability configuration cluster: Assume that cluster nodes are serially upgraded and must be restarted to make upgrades take effect. If the application accesses a node that is being restarted, the node is unavailable.

1. Log on to the ApsaraDB for Cassandra console. Click the name of the cluster whose version is to be upgraded. The Cluster Details page appears.

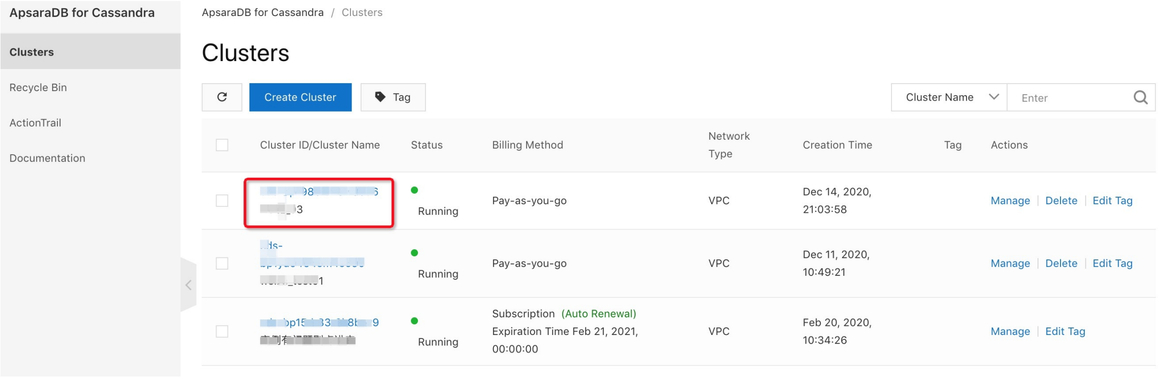

| | Cluster ID/Cluster Name | Status | Billing Method | Network Type | Creation Time | Tag | Actions |
|---|---|---|---|---|---|---|---|
| □ | [redacted] | ● Running | Pay-as-you-go | VPC | Dec 14, 2020, 21:03:58 | | Manage \| Delete \| Edit Tag |
| □ | [redacted] | ● Running | Pay-as-you-go | VPC | Dec 11, 2020, 10:49:21 | | Manage \| Delete \| Edit Tag |
| □ | [redacted] | ● Running | Subscription (Auto Renewal)<br>Expiration Time Feb 21, 2021, 00:00:00 | VPC | Feb 20, 2020, 10:34:26 | | Manage \| Edit Tag |

2. On the **Cluster Details** page, click **Upgrade Minor Version**.

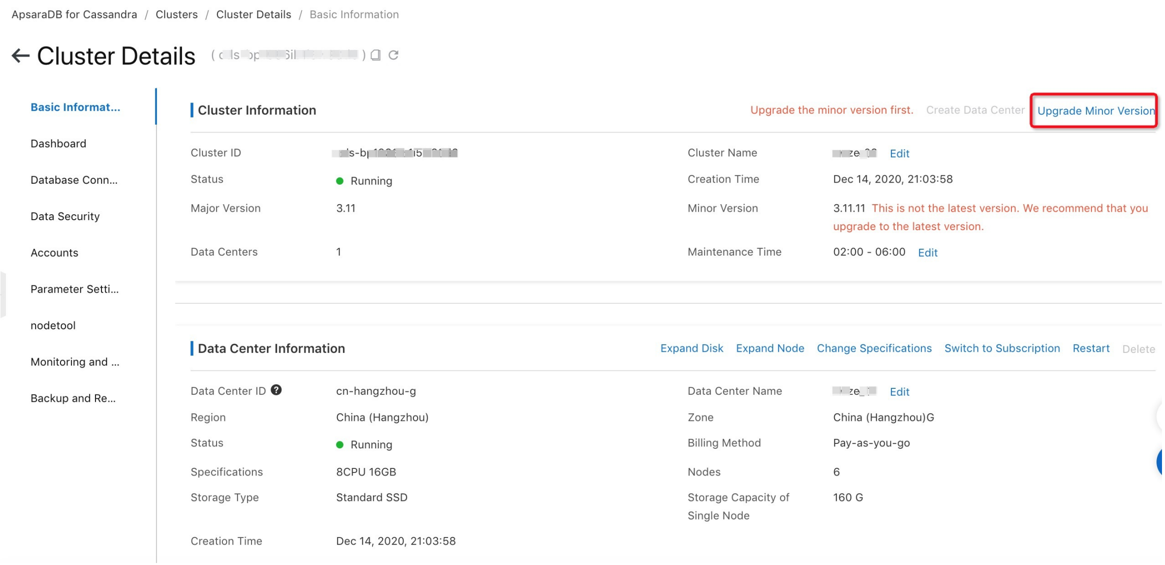

Cluster Information

Upgrade the minor version first.

| | | | |
|---|---|---|---|
| Cluster ID | [redacted] | Cluster Name | [redacted] Edit |
| Status | ● Running | Creation Time | Dec 14, 2020, 21:03:58 |
| Major Version | 3.11 | Minor Version | 3.11.11 This is not the latest version. We recommend that you upgrade to the latest version. |
| Data Centers | 1 | Maintenance Time | 02:00 - 06:00 Edit |

Data Center Information

| | | | |
|---|---|---|---|
| Data Center ID | cn-hangzhou-g | Data Center Name | [redacted] Edit |
| Region | China (Hangzhou) | Zone | China (Hangzhou)G |
| Status | ● Running | Billing Method | Pay-as-you-go |
| Specifications | 8CPU 16GB | Nodes | 6 |
| Storage Type | Standard SSD | Storage Capacity of Single Node | 160 G |
| Creation Time | Dec 14, 2020, 21:03:58 | | |

3. lick **OK** to start the version upgrade.

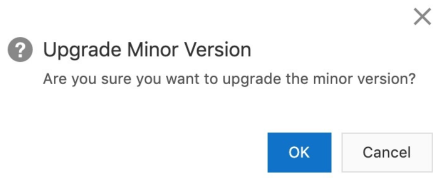

Upgrade Minor Version

Are you sure you want to upgrade the minor version?

# 8.Description of ApsaraDB for Cassandra audit logs

ApsaraDB for Cassandra has been integrated with ActionTrail. This topic describes the ApsaraDB for Cassandra operation logs that are recorded in ActionTrail.

## Background information

ActionTrail is a service that queries and delivers the resource action records of your Alibaba Cloud account. You can view and retrieve behavioral logs in the ActionTrail console. ActionTrail also allows you to deliver logs to Log Service Logstores or a specified Object Storage Service (OSS) bucket. This meets requirements of features such as real-time auditing and problem backtracking and analysis.

## Description of ApsaraDB for Cassandra operation logs

The main content in the ActionTrail logs of ApsaraDB for Cassandra is API call events. For OpenAPI events, the value of the eventType parameter recorded in ActionTrail is `ApiCall`.

## Examples of ApsaraDB for Cassandra operation logs

The following example shows the log of creating an ApsaraDB for Cassandra instance. This log is recorded in ActionTrail. The log records the detailed information about the actions of an ApsaraDB for Cassandra cluster.

```
{
 "eventId": "def79400-0f1e-489a-a1c2-b*********",
 "eventVersion": 1,
 "eventSource": "cassandra.aliyuncs.com",
 "userAgent": "AliyunConsole",
 "eventType": "ConsoleOperation",
 "referencedResources": {
  "ACS::Cassandra::Cluster": [
   "cds-t4n7c886*******"
  ]
 },
 "userIdentity": {
  "accountId": "17926974*****",
  "principalId": "227691078*******",
  "type": "ram-user",
  "userName": "cassandra"
 },
 "serviceName": "Cassandra",
 "requestId": "def79400-0f1e-489a-a1c2-b*******",
 "eventTime": "2020-11-17T14:42:34Z",
 "isGlobal": false,
 "acsRegion": "ap-southeast-1",
 "eventName": "Create"
}
```