ALIBABA CLOUD

# Alibaba Cloud

## ApsaraDB for Cassandra
## Product introduction

Document Version: 20220629

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:** 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 `cd /d C:/window` command to enter the Windows system folder. |
| *Italic* | Italic formatting is used for parameters and variables. | `bae log list --instanceid` *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.What is ApsaraDB for Cassandra?

ApsaraDB for Cassandra is a distributed NoSQL database that is developed based on open source Apache Cassandra and integrated with the Alibaba Cloud Database as a Service (DBaaS) features. ApsaraDB for Cassandra has been in development for over a decade and is based on the distributed architecture of Amazon DynamoDB and data model of Google Cloud Bigtable. ApsaraDB for Cassandra uses a distributed, decentralized, and multi-active architecture, and provides features such as elastic scalability, high availability, fault tolerance, tunable consistency, and SQL-like Cassandra query language (CQL). Apache Cassandra is developed to power services for Internet enterprises and their business. It is the most popular wide table database. Alibaba Cloud launched ApsaraDB for Cassandra in August 2019.

ApsaraDB for Cassandra is preferable for the following reasons:

## Reason 1: ApsaraDB for Cassandra is developed for the Internet business

If you have transaction-related services, we recommend that you use MySQL or PolarDB. MySQL or PolarDB can provide transaction and SQL features to run the business operating software system (BOSS), custom relationship management (CRM) system, and enterprise resource planning (ERP) system. However, for Internet business that requires high availability, high concurrency, massive storage, tunable consistency, and elasticity, we recommend that you use ApsaraDB for Cassandra rather than MySQL.

- High availability: ApsaraDB for Cassandra services are not affected when individual nodes fail. You must configure redundancies for nodes and data replicas and set the read/write consistency level to Quorum.
- High scalability: ApsaraDB for Cassandra can provide storage from 160 GB to 10 PB and up to tens of millions of queries per second (QPS). You can add more nodes and data centers to your services as needed.
- Tunable consistency: ApsaraDB for Cassandra provides high availability and strong consistency by setting the read/write consistency level to Quorum. You can configure low consistency for Internet of Things (IoT) services based on your business features to achieve higher performance at a lower cost.

We recommend that you use ApsaraDB for Cassandra for services such as online chatting and messaging, news feeds, history orders, object storage, shopping carts, billing systems, and comments. The services also include data-driven services such as risk control, recommendation, user profiling, IoT, and log analysis.

## Reason 2: ApsaraDB for Cassandra is easy-to-use for developers and DBAs

A MySQL developer or database administrator (DBA) can master the skills of ApsaraDB for Cassandra databases in a short period of time. ApsaraDB for Cassandra is easy to use and manage and provides the following benefits:

- SQL-like CQL enables MySQL database developers and other traditional database developers to directly use ApsaraDB for Cassandra databases.
- ApsaraDB for Cassandra provides a security authentication system, a comprehensive authentication system, and an SSL certificate mechanism for DBAs.

- ApsaraDB for Cassandra provides complete index services such as LocalIndex, materialized views (MVs), and SSTable-Attached Secondary Index (SASI) for developers.
- Native Apache Cassandra provides client drivers (non-Thrift mode) for multiple programming languages, such as Java, Python, PHP, .NET, and Node.js. The performance of native client drivers is consistent with that of Java clients.

## Reason 3: ApsaraDB for Cassandra works with X-Pack Spark to facilitate big data storage

Apache Spark is a mainstream default engine for big data processing and Apache Cassandra is a mainstream NoSQL big data database. The combination of Apache Cassandra and Apache Spark allows you to manage data-driven services such as risk control, recommendation, and IoT. ApsaraDB for Cassandra provides the X-Pack Spark engine to facilitate your usage.

## Reason 4: 24/7 support is available to ensure stability and security

ApsaraDB for Cassandra provides enterprise-grade disaster recovery features such as multiple data centers deployed in a zone and backup and restoration. Alibaba Cloud database team offers 24/7 support to ensure database stability and security and provides enhanced performance and service-related features for your database.

# 2.Benefits

| Item | ApsaraDB for Cassandra | User-created database |
| --- | --- | --- |
| Service availability | 99.9% | Requires you to guarantee your own service availability. |
| Backup and restoration | Provides the backup and restoration feature for clusters. | Not provided. |
| Maintenance cost | No O&M required. | Requires professional database administrators (DBAs) and high O&M costs. |
| Deployment and scaling | Supports real-time activation, fast deployment, and elastic scaling. | Requires hardware procurement, hosting of data centers, and machine deployment, which is time-consuming. |
| Multi-active geo-redundancy | Deployment can be completed in one click. | Deployment is complex and takes a long period of time. |
| Kernel management | Supports automatic upgrades to fix vulnerabilities at the earliest opportunity. It saves you the efforts to manage kernel versions. Parameter settings are optimized to maximize the utilization of system resources. ApsaraDB for Cassandra kernel development team fixes bugs and provides new features. | Requires you to backport community patches and optimize parameters. |
| Performance optimization | Throughput is increased by 20% and can be increased by up to 100%. | Not provided. |
| Analysis | Seamless integrated with X-Pack Spark to meet requirements for streaming and analysis. | Not provided. |

# 3.System architecture

The following architecture shows an ApsaraDB for Cassandra instance that contains two data centers. Each data center contains four nodes.

## Basic architecture:

- An ApsaraDB for Cassandra instance can contain one to three data centers. Each data center contains one to 500 nodes. A Cassandra process is deployed on each node.
- ApsaraDB for Cassandra instances do not adopt the master-proxy node architecture. All the nodes are peers. Logically, the system specifies two nodes in each data center as the seed nodes. However, seed nodes do not require process deployment because seed nodes bring little impact on the system.
- The client connects to each server node to ensure high performance.
- The number of replicas of each data center can be customized. Replicas can be used for data center disaster recovery, backup, and offline analysis.

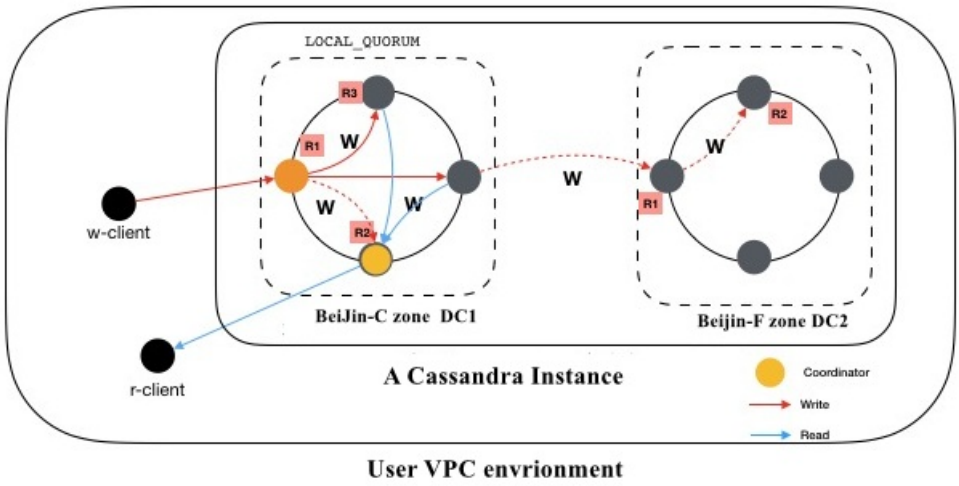


## Nodes

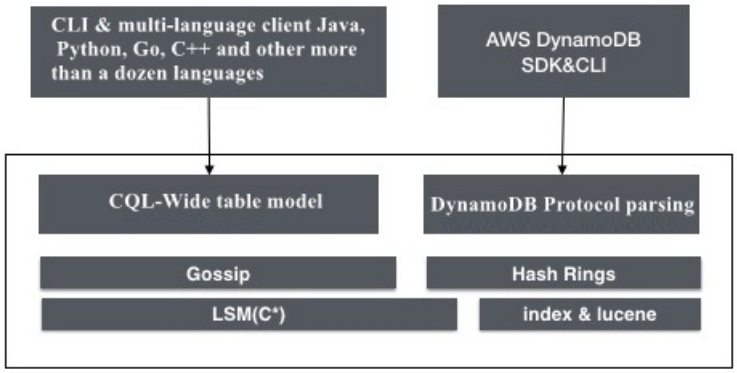


The storage space of a node ranges from 160 GB to 16 TB. The storage space of 500 nodes is 8000 TB that is about 8 PB.

# 4.Scenarios

## Internet services

ApsaraDB for Cassandra supports access requests that require high concurrency and low latency, and provides high availability and elastic scaling features. ApsaraDB for Cassandra is applicable to online Internet scenarios that involve a large amount of data, such as logs, messages, feed streams, orders, bills, and websites.

## Multi-active geo-redundancy

Apache Cassandra supports deployment in multiple data centers to ensure high availability and disaster recovery. ApsaraDB for Cassandra allows you to add more data centers and configure multiple replicas for data center disaster recovery, backup, and offline analysis.

## Flexible services

ApsaraDB for Cassandra provides flexible data models to easily modify table schemas. It is suitable for startup business that is ever-changing.

## Write-intensive, statistical, and analytical operations

ApsaraDB for Cassandra provides optimized write throughput to support high write performance in multiple client threads and handle burst peaks. This feature makes ApsaraDB for Cassandra ideal for scenarios that require significantly more writes than reads, such as user status update, social networking, recommendations and comments, and application statistics.

## Data-driven services

ApsaraDB for Cassandra supports hundreds of nodes in a cluster to store a large amount of data. In user behavior analysis scenarios that involve a large amount of data, ApsaraDB for Cassandra can store user behavior data and user profiles created from multiple data sources to provide big data risk control and recommendation services.

# 5.Terms

| Term | Description |
| --- | --- |
| Cassandra | ApsaraDB for Cassandra, an online distributed NoSQL database. It supports the SQL-like Cassandra query language (CQL) and multi-active geo-redundancy to provide enterprise-grade features such as security assurance, disaster recovery, monitoring, and backup and restoration. |
| CQL | The Cassandra query language. It is a SQL-like query language provided by Cassandra. For more information, visit The Cassandra Query Language (CQL). |
| data center | A data center is a node group that has independent power supply and network in a region. Multiple data centers within an instance are connected through the internal network and have little network latency. Faults are isolated between data centers. |
| partitioner | A partitioner determines the policy of distributing data evenly among nodes. ApsaraDB for Cassandra uses Murmur3Partitioner by default. |
| number of replicas | The number of replicas indicates the number of data copies in a cluster. For example, if the number of replicas is two, each row in a cluster has two replicas that are on different nodes. You can specify the number of replicas when you create a keyspace. |
| replica policy | A replica policy determines the node on which replicas are stored. You can specify the replica policy when you create a keyspace. We recommend that you use NetworkTopologyStrategy to deploy a cluster in multiple data centers. |
| keyspace | A keyspace contains multiple tables. You can specify replica policies for keyspaces. |

# 6.Availability

ApsaraDB for Cassandra of the high availability configuration does not require timeout detection and log playback, causing minimum impact (measured in milliseconds) on the system.

| | QUORUM | High availability mechanism | Raft | Region |
|---|---|---|---|---|
| Fault detection | Coordinator node and client retry, no need for timeout detection, log playback, or leader reelection | Client retry, timeout detection, and log playback | Client retry and leader reelection | Client retry, timeout detection, and log playback |
| Response time | milliseconds | 10 seconds to 10 minutes | 10 seconds to 10 minutes | 1 minute to 10 minutes |
| Typical system | Cassandra | RDS, MongoDB, and Redis of the primary/secondary architecture | N/A | HBase |

## Recommended configuration

The basic requirements for high availability configurations are as follows:

## Fault testing

Kill one node out of three nodes that each has 2 cores CPU and 4 GB of memory when intensive write operations are performed. The client operation status is as follows:

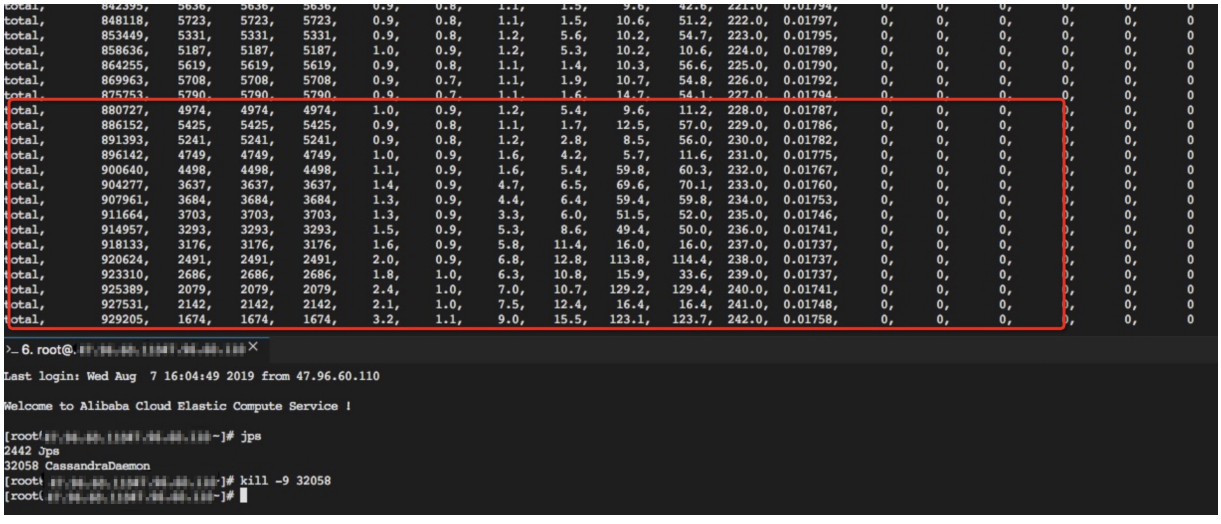


The Cassandra service is not affected when a node is killed.

# 7.Expert service

If you want to learn more about ApsaraDB for Cassandra or have any questions about it, use the following methods:

1. Visit the ApsaraDB for Cassandra developer community.
2. Scan the following QR code to join the Cassandra+Spark community DingTalk group for expert service.

# 8.Scalability

ApsaraDB for Cassandra features horizontal scalability. A single data center can contain up to 500 nodes. An instance can contain up to three data centers.

- Contains no master or proxy nodes to save resources.
- Supports queries per second (QPS) scaling from 1,000 to 100 million when more nodes are added to avoid central bottleneck.
- Supports scaling up nodes. The basic specification is 1 core CPU and 2 GB of memory.
- Supports storing large amounts of data from 160 GB to 10 PB.
- Supports adding more data centers.

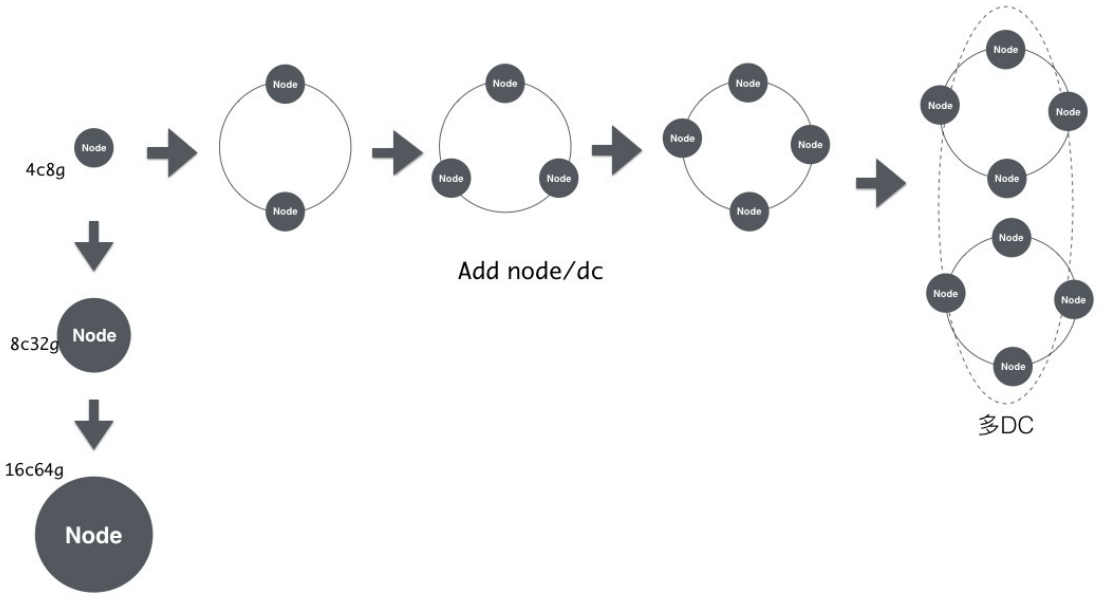


Scalability

# 9.Release notes

The major version is 3.11. The minor version is released as follows:

- 3.11.5 (release date: August 8. 2019)
  - The features of the Apache Cassandra community edition 3.11.3 are supported.
  - Bugs are fixed.
  - Supports public and internal access.
  - The root permissions are supported. The delete permissions on system tables are denied to avoid accidental system table deletion.