

Alibaba Cloud ApsaraDB for POLARDB

Product Introduction

Issue: 20190909

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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, 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 content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content 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 contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information 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.
	This warning information 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 restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK.
<code>Courier font</code>	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>switch {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 What is ApsaraDB for POLARDB?.....	1
2 Benefits.....	5
3 Architecture.....	7
4 Limits.....	10

1 What is ApsaraDB for POLARDB?

ApsaraDB for POLARDB is a next-generation relational database service developed by Alibaba Cloud. It is compatible with MySQL, PostgreSQL, and Oracle databases. With superior performance in storage and computing, ApsaraDB for POLARDB can meet the diverse requirements of enterprises. An ApsaraDB for POLARDB cluster has a maximum storage space of 100 TB and can be configured with a maximum of 16 nodes.

ApsaraDB for POLARDB uses a storage and computing-separated architecture, in which all compute nodes share one copy of data. It can achieve vertical scaling within minutes and crash recovery within seconds. It ensures global data consistency, and offers free services for data backup and disaster recovery. ApsaraDB for POLARDB integrates the benefits of commercial databases and open source cloud databases. Commercial databases are stable, reliable, high-performance, and scalable while open source databases are easy to use and self-iterative. For example, the time a POLARDB for MySQL database takes to return results for a query reduces by five times than that a MySQL database takes. However, the cost of a POLARDB for MySQL database is only 10% that of a MySQL database.

Cluster architecture with separate storage and computing: ApsaraDB for POLARDB adopts a cluster architecture. Each cluster contains one read/write node (primary node) and multiple read-only nodes. All nodes share the data in a Polar store by using a distributed Polar file system.

Read/write splitting: POLARDB for MySQL uses a built-in proxy to handle external requests. When apps use a cluster address, the proxy handles all requests sent from the apps before forwarding the requests to nodes. You can use the proxy for authentication and protection and use it to achieve automatic read/write splitting. The proxy can parse SQL statements, send write requests (such as transactions, UPDATE, INSERT, DELETE, and DDL operations) to the primary node, and distribute read requests (such as SELECT operations) to multiple read-only nodes. With the proxy, apps can access POLARDB for MySQL as easily as they access a single-node MySQL database. The proxy only supports POLARDB for MySQL. We are working on support for POLARDB for PostgreSQL and POLARDB for Oracle.

Terms

Familiarize yourself with the following terms to gain a better understanding of ApsaraDB for POLARDB. This helps you to find optimal purchase strategies and use the ApsaraDB for POLARDB service based on your needs.

- **Cluster:** ApsaraDB for POLARDB adopts a cluster architecture. Each cluster contains one primary node and multiple read-only nodes.
- **Region:** specifies the region in which a data center resides. You can achieve optimal read/write performance if ApsaraDB for POLARDB clusters and ECS instances are located in the same region.
- **Zone:** A zone is a distinct location that operates on independent power grids and networks within a region. All zones in a region provide the same services.
- **Specification:** specifies the resources configured for each node, such as 2 CPU cores and 4 GB.

Benefits

ApsaraDB for POLARDB is compatible with MySQL, PostgreSQL, and Oracle databases. ApsaraDB for POLARDB has the following benefits:

- **Large storage space**

A maximum storage space of 100 TB for an ApsaraDB for POLARDB cluster overcomes the limit of a single host and alleviates the need to purchase multiple instances for database sharding. ApsaraDB for POLARDB simplifies application development and reduces O&M workloads.

- **Cost-effectiveness**

When you add a read-only node in an ApsaraDB for POLARDB cluster, you only need to pay for computing resources because of storage and computing separation. In contrast, traditional databases charge you for both computing and storage resources.

- **Elastic scaling within minutes**

You can quickly scale up an ApsaraDB for POLARDB cluster because of storage and computing separation as well as shared storage.

- Read consistency

The cluster address uses log sequence numbers (LSNs) to ensure global consistency in reading data and to avoid inconsistency caused by synchronization latency between the primary node and read-only nodes.

- Millisecond-level latency in physical replication

ApsaraDB for POLARDB uses redo log-based physical replication from the primary node to read-only nodes instead of binlog-based logical replication to improve the efficiency and stability. No latency is incurred for databases even if you perform DDL operations for a large table, such as adding indexes or fields.

- Unlocked backup

You can back up a database with a size of 2 TB within 60 seconds by using snapshots. Backup can be performed at any time on a day without any impacts on apps. During the backup process, the database will not be locked.

Pricing

For more information about pricing, see [#unique_4](#).

Use ApsaraDB for POLARDB

You can use the following methods to manage ApsaraDB for POLARDB clusters, for example, to create clusters, databases, and accounts:

- **Console:** provides a visualized and easy-to-use Web interface.
- **CLI:** All operations available in the console can be performed by using the command-line interface (CLI).
- **SDK:** All operations available in the console can be performed by using the SDK.
- **API:** All operations available in the console can be performed by using API operations.

After creating an ApsaraDB for POLARDB cluster, you can connect to the cluster by using the following methods:

- **DMS:** You can [connect to an ApsaraDB for POLARDB cluster by using Data Management \(DMS\)](#) and develop databases on the Web interface.
- **Client:** You can connect to an ApsaraDB for POLARDB cluster by using a database client, such as MySQL-Front and pgAdmin.

Related services

- **ECS:** When Elastic Compute Service (ECS) instances access ApsaraDB for POLARDB clusters in the same region, the optimal performance of ApsaraDB for POLARDB clusters is achieved. ECS instances and ApsaraDB for POLARDB clusters compose a typical business architecture.
- **ApsaraDB for Redis:** ApsaraDB for Redis provides database services that use hybrid storage of memory and hard disks to ensure data consistency. You can combine ECS instances, ApsaraDB for POLARDB clusters, and ApsaraDB for Redis instances to handle a large number of read requests and reduce the response time.
- **ApsaraDB for MongoDB:** ApsaraDB for MongoDB provides a stable, reliable, and scalable database service that is compatible with the MongoDB Wire Protocol. To meet diverse business demands, you can store structured data in ApsaraDB for POLARDB and store unstructured data in ApsaraDB for MongoDB.
- **DTS:** You can use Data Transmission Service (DTS) to migrate on-premises databases to ApsaraDB for POLARDB.
- **OSS:** Object Storage Service (OSS) is a cloud storage service that features significant storage capacity, security, cost-effectiveness, and reliability.

2 Benefits

Easy to use

POLARDB is fully compatible with MySQL, which means that the code, applications, and drivers you use with your existing MySQL databases can be used with POLARDB with little or no changes required.

High performance

- The POLARDB architecture improves on the MySQL kernel, and adopts physical replication, RDMA protocol, and shared distributed storage, improving the read performance up to six times when compared to MySQL.
- A POLARDB cluster contains one primary instance and up to 15 read-only instances (with at least one read-only instance acting to provide active-active high availability solutions) so that you can perform concurrent read and write operations.

Massive storage

Distributed block storage and file system is used in the POLARDB architecture to support a capacity of up to 100 TB per database instance.

High availability

POLARDB performs automated failovers targeted at the primary instance, or a read-only instance, by using an active-active high availability architecture. This architecture brings better system access performance than a traditional active-standby setup at the same cost.

Data security and reliability

POLARDB provides various security measures for your database access, storage, and management. These include managing data access by whitelisting IP addresses, network isolation using Virtual Private Cloud (VPC), encryption of data in transit using SSL, and storage of data in multiple copies.

Low cost


POLARDB lowers storage costs by applying a shared storage mechanism.

Scalability and expandability

POLARDB enables easy CPU and memory expansion by using container virtualization and shared distributed block storage technology. It takes only 2-3 minutes to add a secondary instance to your database. Based on your database usage, the database storage will automatically scale towards the upper limit of your storage (depending on the specifications of your instance), with no impact on business continuity.

3 Architecture

The architecture of the cloud-native POLARDB is shown in the following figure.



One primary instance, multiple read-only instances

A POLARDB cluster contains one primary instance and up to 15 read-only instances (with at least one read-only instance to provide active-active high availability support). The primary instance processes read and write requests, while read-only instances process read requests only. POLARDB provides highly available services by performing active-active failover targeted at the primary instance or a read-only instance.

Compute and storage separation

POLARDB separates compute processes from storage processes, allowing the database capacity to scale up and down to meet your application needs in Alibaba Cloud. DB servers are used only to store metadata, while data files and redo logs are stored in remote chunk servers. DB Servers only need to synchronize redo log metadata between each other, which significantly lowers the data latency between the primary instance and read-only instances. If the primary instance fails, a read-only instance can be rapidly promoted to be primary.

Read/write splitting

Read/write splitting is enabled for POLARDB clusters by default, providing transparent, highly available, and self-adaptive load balancing for your database. The read/write splitting feature automatically routes requests directed at the read/write splitting connection string. It passes write requests to the primary instance and passes read requests to either the primary instance or a read-only instance based on the load of each instance. This allows the database to handle large numbers of concurrent requests.

High speed network connection

To ensure strong I/O performance, high speed network connection is enabled between the DB server and the chunk server, and data are transferred using the Remote Direct Memory Access (RDMA) protocol.

Shared distributed storage

Sharing the same group of data copies among multiple DB servers, rather than storing a separate copy of data for each DB server, significantly reduces your storage cost. The distributed storage and file system allows automatically scaling up database

storage capacity, regardless of the storage capacity of each single database server.

This enables your database to handle up to 100 TB of data.

Multiple data replicas, Parallel-Raft protocol

Chunk servers maintain multiple data replicas to ensure reliability, and comply with the Parallel-Raft protocol to guarantee consistency among these replicas.

4 Limits

Specification limits

Specifications	Memory	Memory occupied by the file system	Maximum storage	Maximum number of files
polar.mysql.x2.medium	4 GB	450 MB	5,000 GB	1,026,048
polar.mysql.x8.medium	16 GB	850 MB	1,000 GB	2,050,048
polar.mysql.x8.large	32 GB			
polar.mysql.x8.xlarge	64 GB			
polar.mysql.x8.2xlarge	128 GB			
polar.mysql.x8.4xlarge	220 GB			

Note:

- Maximum number of files: includes user files, database system files (approximately 100), and log files (run the SHOW POLAR LOGS command to view the number of log files). A POLARDB table occupies two files. A partition table occupies N+2 files (N indicates the number of partitions). The following error message will pop up if you try to create new tables after reaching the upper limit of files:

ERROR 3017 (HY000): Too many files. PolarDB only supports 2048 files every 10GB disk size. Please drop some tables/databases before creating new tables

Specifically, you must delete some tables or upgrade the specifications of your database instance.

- Memory occupied by the file system: The memory occupied by the file system when your database reaches the upper limit of storage and is able to read and write normally (that is, neither when the instance is under performance stress testing, nor when you perform DDL on large tables). The file system occupies fewer

memory resources than specified in the table if your instance has not reached its storage limit.

Other limits

- The length of table name: The table name (or file name) in any database instance cannot exceed 63 English letters in length.
- You are unable to set the serializable isolation level for your instances.
- You cannot manage accounts in the console because your account is an advanced user by default. The permissions of POLARDB accounts are similar to those of RDS accounts for MySQL as advanced users.