Alibaba Cloud ApsaraDB for POLARDB

User Guide for MySQL

Issue: 20190918

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.
`Courier font`	It is used for commands.	Run the `cd / d C :/ windows` command to enter the Windows system folder.
Italics	It is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	It indicates that it is a optional value, and only one item can be selected.	`ipconfig` *[-all|-t]*

Style	Description	Example
{} or {a|b}	It indicates that it is a required value, and only one item can be selected.	`swich {stand | slave}`

Contents

1 Overview

POLARDB is a next-generation cloud-based service developed by Alibaba Cloud for relational databases, which is fully compatible with MySQL. It is developed based on a distributed storage architecture and delivers up to six times throughput of MySQL. POLARDB provides high-capacity, low-latency online transaction processing (OLTP) services. It also offers cost-effective and scalable services.

Basic concepts

- Cluster: A POLARDB cluster contains one primary instance and a maximum of 15 read-only instances (at least one, in order to provide active-active high availability support). A POLARDB cluster ID starts with pc, which stands for POLARDB cluster.
- Instance: An instance is an independent database server in which you can create and manage multiple databases. An instance ID starts with pi, which stands for POLARDB instance.
- Database: A database is a logical unit created in an instance. The name of each POLARDB database under the same instance must be unique.
- Region and zone: Each region is a separate geographic area. Zones are distinct locations within a region that operate on independent power grids and networks. For more information, see Alibaba Cloud's Global Infrastructure.

Console

Alibaba Cloud offers a web-based and easy-to-use console where you can manage various products and services including POLARDB. In the console, you can create, access, and configure your POLARDB database.

For more information about the console layout, see Alibaba Cloud console.

The logon page for the POLARADB console is https://polardb.console.aliyun.com.

2 Read/write splitting

POLARDB for MySQL clusters support read/write splitting. To enable read/write splitting, you must use the cluster address when sending requests from apps to a POLARDB for MySQL cluster. Then, the built-in proxy of the POLARDB for MySQL cluster forwards write requests to the primary node, and forwards read requests to the primary node or a read-only node based on the loads. The loads for a node is indicated by the number of requests that the node is handling.

Benefits

- Read consistency

  ApsaraDB for POLARDB uses a proxy to achieve read/write splitting, load balancing , and read consistency. This topic details how ApsaraDB for POLARDB achieves read consistency without incurring much workloads for a primary node. The proxy tracks the log sequence number (LSN) of the redo log for each node. Each time a log stored in the primary node is updated, the new LSN of the log is labeled as the session LSN. If a new read request arrives within a session, the proxy compares the session LSN and the LSN of the log stored in each node. Then, the proxy forwards the request to a read-only node where the LSN is equal to or greater than the session LSN. To enable an instant replication, after the primary node handles a write request, it returns the result to the client and replicates data to read-only nodes at the same time. Physical replication is rapid, which allows read-only nodes to update data before subsequent read requests within the session arrive.

- Built-in proxy for read/write splitting

  You can use a self-created proxy on the cloud to achieve read/write splitting. However, excessive latency may occur because data is parsed and forwarded by multiple components before arriving at a database. ApsaraDB for POLARDB uses a built-in proxy to reduce latency and enhance query performance. The proxy is deployed within the existing secure links with no additional components that consume much time in parsing and forwarding requests.

- O&M efficiency

  For traditional database services, you need to spend much effort to achieve read/write splitting. You must configure the addresses of all the primary node and read-only nodes in apps. Then, you need to distribute write requests to the primary node and read requests to read-only nodes.

  ApsaraDB for POLARDB provides an optimal solution by using a cluster address. After apps establish connections to a cluster by using the cluster address, read and write requests are automatically forwarded to the intended primary node and read-only nodes. You can view the forwarding results in the console. The cluster address minimizes your effort and cost in maintaining the ApsaraDB for POLARDB cluster.

  You can expand the capacity of an ApsaraDB for POLARDB cluster by adding read-only nodes, which alleviates your need to make any modifications for apps.
- Health checks for nodes

  ApsaraDB for POLARDB automatically performs health checks for all nodes in a cluster. If a node shuts down or has a latency that exceeds the threshold, requests will not be distributed to the node. This ensures that apps can access the ApsaraDB for POLARDB cluster even if a node crashes. After the node recovers, ApsaraDB for POLARDB automatically adds it into the list of nodes that are available to apps.
- Free service

  The read/write splitting feature is available for free.

Forwarding logic

- Forwarding logic in read/write splitting mode
  - The following requests are forwarded to the primary node:
    - All DML operations, such as INSERT, UPDATE, and DELETE operations
    - All DDL operations, such as creating databases or tables, deleting databases or tables, and changing table schemas or permissions
    - All requests in transactions
    - Queries by using user-defined functions
    - Queries by using stored procedures
    - EXECUTE statements
    - Multi-statements
    - Requests that involve temporary tables
    - SELECT last_insert_id()
    - All requests to query or modify user environment variables
    - SHOW PROCESSLIST statements
    - KILL statements in SQL (not KILL commands in Linux)
  - The following requests are forwarded to the primary node or read-only nodes:
    - Non-transactional read requests
    - COM_STMT_EXECUTE commands
  - The following requests are forwarded to all nodes:
    - All requests to modify system environment variables
    - USE commands
    - COM_STMT_PREPARE commands
    - Commands such as COM_CHANGE_USER, COM_QUIT, and COM_SET_OPTION
- Forwarding logic in read-only mode:
  - DDL and DML operations are not allowed.
  - Requests are forwarded to read-only nodes based on load balancing.
  - Requests cannot be forwarded to the primary node.

Limits

- If you run a multi-statement or call a stored procedure, all subsequent requests during the current session are forwarded to the primary node. To re-activate the

read/write splitting feature, you must disable the current connection and establish a new connection.

- In read-only mode, you cannot use the cluster address to configure environment variables. For example, if you run `set @endtime = now(); select * from tab where dt < @endtime`, you may not retrieve the intended environment variable.
- When you use views, session consistency cannot be guaranteed. For example, if you run `CREATE VIEW tab2 AS SELECT * FROM tab1; INSERT INTO tab1 (key, value) (1, 1); SELECT * FROM tab2 where key = 1;`, the result may not be returned because of timeout.

Apply for or modify a cluster address

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster, and click the cluster ID in the `Cluster Name` column.
4. On the Overview page, view Cluster Endpoints (Recommended) in the Connection Information section.
5. Click Apply, and in the dialog box that appears, click Confirm. Then, you can see the cluster address on the Overview page.

Note:

If an existing cluster does not have a cluster address, you must manually apply for the cluster address. The cluster address is automatically assigned to the new purchased cluster. If an ApsaraDB for POLARDB cluster has the cluster address, you can perform the following step to modify the address.

6. Click Modify. In the Modify Endpoint dialog box, enter a new cluster address and click Submit.

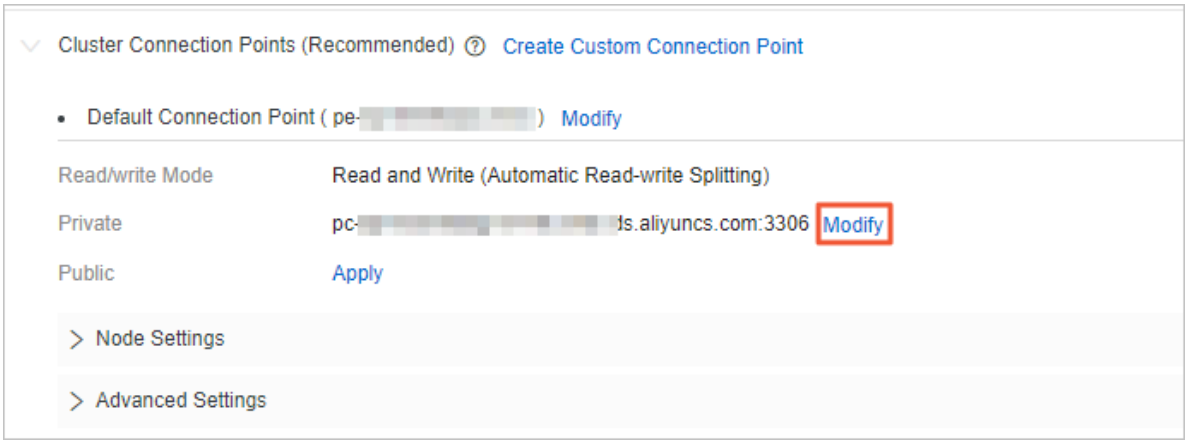

FAQ

1. Why cannot I retrieve a newly inserted record?

   In a read/write splitting architecture, latency may occur during data replication from the primary node to read-only nodes. ApsaraDB for POLARDB ensures that the updates within a session can be retrieved.

2. Why do read-only nodes have no workloads?

   Requests in transactions are forwarded to the primary node by default. If you use sysbench to benchmark an ApsaraDB for POLARDB cluster, you can add --oltp-skip-trx=on (for sysbench 0.5) or --skip-trx=on (for sysbench 1.0) in the code. This ensures that read requests in transactions can be forwarded to read-only nodes. If a large number of transactions incur too much workloads for the primary node, you can submit a ticket to enable distribution of transactions to read-only nodes.

3. Why does a node receive more requests than other nodes?

   Requests are distributed to each node based on node workloads. The node that has low workloads will receive more requests.

4. Can I retrieve the query result with no latency?

   Under common workloads, an ApsaraDB for POLARDB cluster can transmit data from the primary node to read-only nodes with a latency of milliseconds. If you want to retrieve the query result with no latency, you can use the primary address to manually send all requests to the primary node.

5. Will new read-only nodes be automatically available to receive read requests?

   If a read-only node is added, the proxy will forward requests to the node after new connections between apps and an ApsaraDB for POLARDB cluster are established. You need to disable the current connection and establish a new connection. For example, you can restart apps.

Related API operations

Operation	Description
#unique_5	Creates a public network address for an ApsaraDB for POLARDB cluster.
#unique_6	Creates a custom cluster address for an ApsaraDB for POLARDB cluster.
#unique_7	Queries the address information of an ApsaraDB for POLARDB cluster.
#unique_8	Modifies the configurations of the cluster address for an ApsaraDB for POLARDB cluster.
#unique_9	Modifies the prefix of a default cluster address for an ApsaraDB for POLARDB cluster.
#unique_10	Deletes the cluster address of an ApsaraDB for POLARDB cluster (excluding the private network address).
#unique_11	Deletes the custom cluster address of an ApsaraDB for POLARDB cluster.

3 Pending events

When an ApsaraDB for POLARDB event is pending for processing, you will be notified to handle the event in a timely manner in the console.

For ApsaraDB for POLARDB O&M events, including database software upgrade events and hardware maintenance and upgrade events, you are notified not only by SMS messages, phone calls, emails, or internal messages, but also in the console. You can view the details of each event, including the event type, task ID, cluster name, and switch time. You can also change the switch time.

Prerequisites

There are unprocessed O&M events.

> **Note:**
> If there are unprocessed O&M events, you can see notification badges on the Pending Events page.

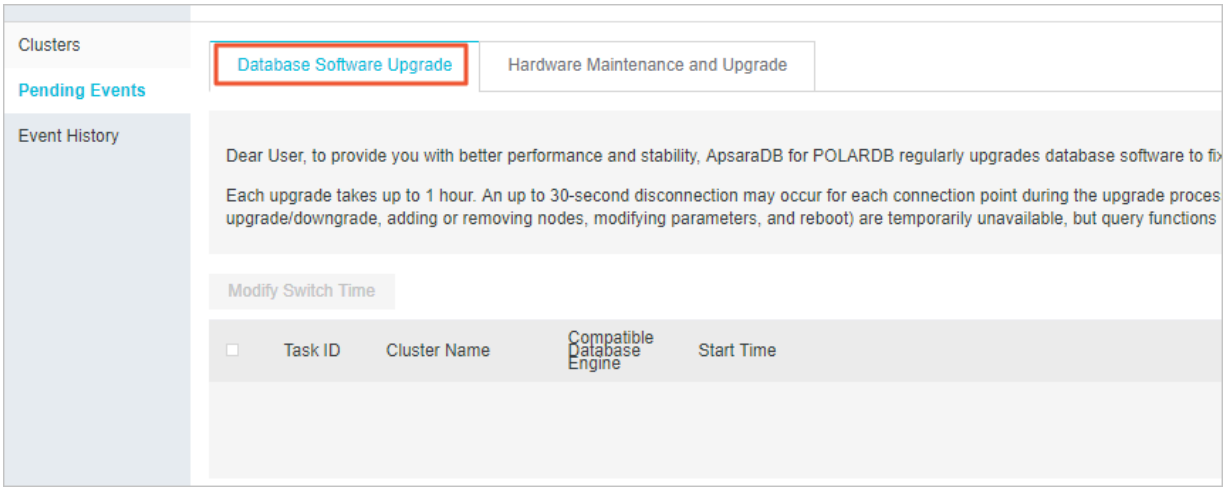

Change the switch time

1. Log on to the ApsaraDB for POLARDB console.
2. In the left-side navigation pane, click Pending Events.

> **Note:**
> For an O&M event for which you must reserve the switch time, a dialog box appears, asking you to complete the reservation as soon as possible.

3. On the Pending Events page, select the type of event that you want to handle.

> **Note:**
> Different notices are displayed on the tabs for different types of events.

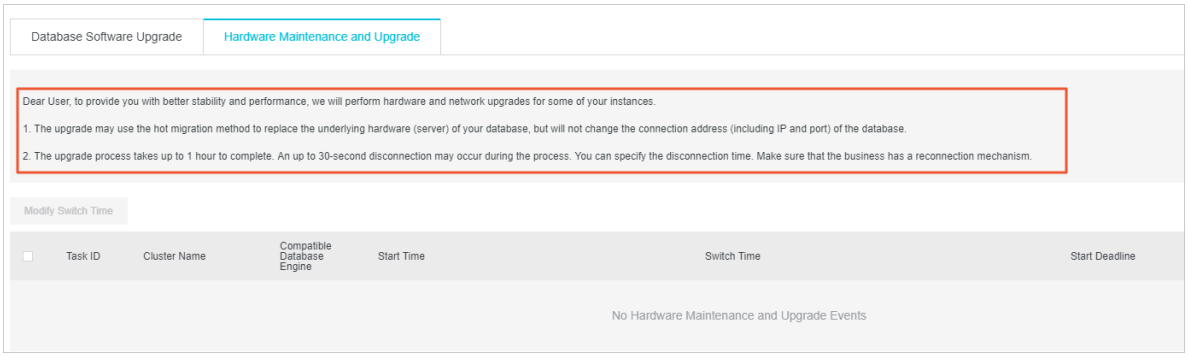

4. View event details in the event list. To change the switch time, select an event, and then click Change Switch Time. In the dialog box that appears, set the switch time, and then click OK.

> **Note:**
> The switch time cannot be later than the latest operation time allowed.

Historical events

You can view completed events on the Event History page.

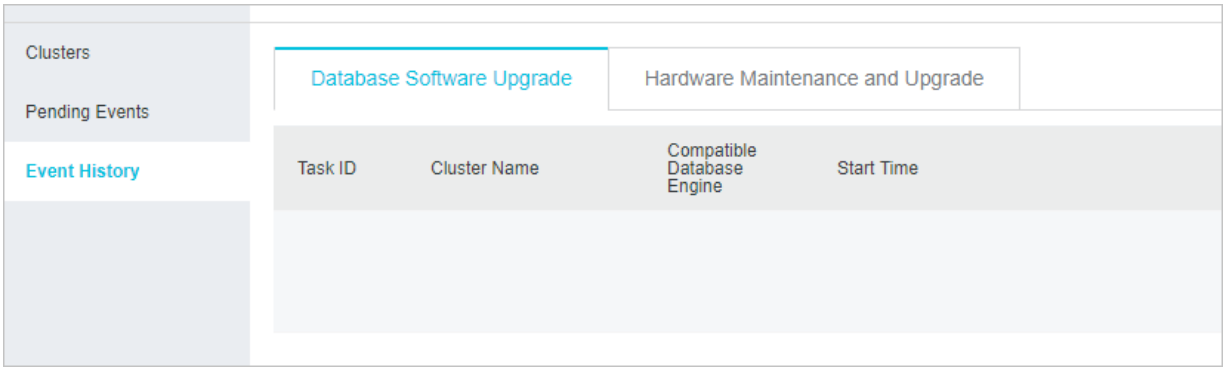

4 Configure a whitelist for a POLARDB for MySQL cluster

After a POLARDB cluster is created, you must add the IP addresses for accessing the cluster to a whitelist, and create an initial account. Otherwise, you cannot access or use the cluster.

Only IP addresses in the whitelist can be used to visit the nodes in the cluster. The whitelist can only be configured on the cluster details page, and is applicable to all nodes in the cluster.

Precautions

- By default, the whitelist only contains an IP address 127.0.0.1. It means that no IP addresses can access the cluster.
- If the whitelist is set to % or 0.0.0.0/0, it allows all IP addresses to access the cluster. This configuration substantially compromises the database security. Therefore, this configuration is not recommended.
- POLARDB cannot automatically obtain the private IP addresses of ECS instances in the VPC. You must add the private IP addresses to the whitelist manually.

Procedure

1. Log on to the POLARDB console.
2. Find the target POLARDB cluster and click its ID.
3. In the Access Information section of the Basics page, click Configure below Whitelists.

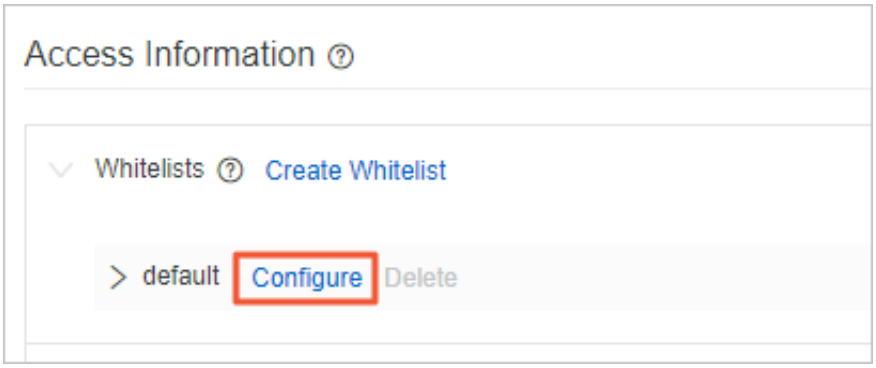

4. In the Configure Whitelist dialog box, add IP addresses to allow these IP addresses to access the POLARDB cluster.
5. Click Submit.

APIs

API	Description
#unique_14	Used to list the IP address that are allowed to access the a POLARDB cluster.
#unique_15	Used to modify the list of IP addresses that are allowed to access the POLARDB cluster.

5 Billing management

5.1 Change the billing method from pay-as-you-go (hourly rate) to subscription

You can change the billing method of a cluster from pay-as-you-go (hourly rate) to subscription to meet your needs. Changing the billing method will not impact the performance of an ApsaraDB for POLARDB cluster.

> Note:
> If a cluster uses a specification that is no longer available, you cannot change the billing method of the cluster to subscription. In this case, you need to #unique_18 before changing the billing method.

Precautions

You cannot change the billing method of a cluster from subscription to pay-as-you-go (hourly rate). Exercise caution before you change the billing method to subscription.

Prerequisites

- The status of your cluster must be Running.
- There are no pending orders for changing the billing method from pay-as-you-go (hourly rate) to subscription. If there are any pending orders, you must pay or discard them on the Orders page.

Procedure

1. Log on to the ApsaraDB for POLARDB console.
2. Select the region where the cluster is located.

3. Find the target cluster. In the Actions column of the cluster, click the More icon, and then select Switch to Subscription.

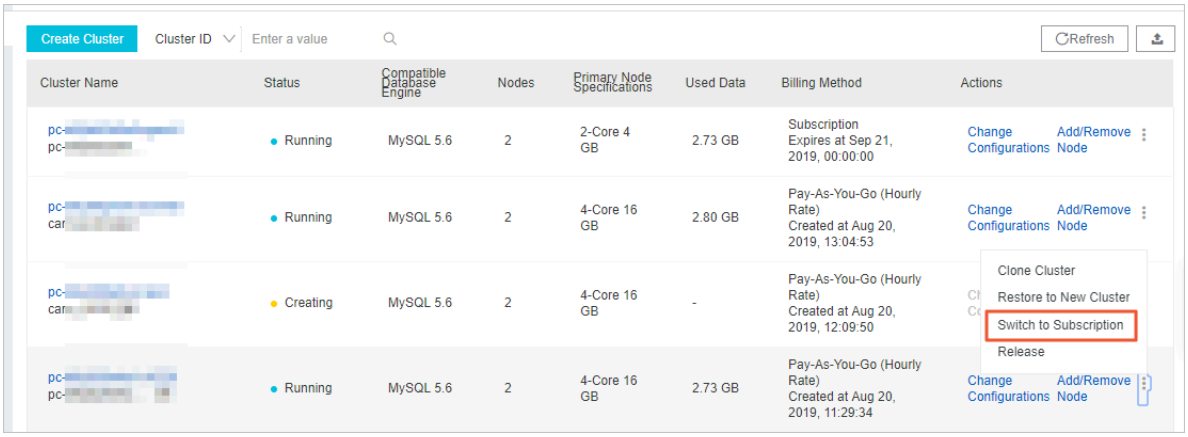

4. Specify the renewal duration, read the ApsaraDB for POLARDB Subscription Agreement of Service, select the check box to agree to it, and then click Activate.

> Note:
>
> - The new billing method will take effect after you complete the payment.
> - If the order is unpaid or payment is unsuccessful, an incomplete order will be listed on the Orders page, and you cannot purchase any new cluster or change the billing method to subscription. You must pay or discard the order before placing a new one.

5.2 Manually renew the subscription to a cluster

You can renew your subscription to clusters in the ApsaraDB for POLARDB console or in the Renew console. In the Renew console, you can renew your subscription to multiple clusters at the same time.

> Note:
> Clusters purchased through the pay-as-you-go (hourly rate) billing method do not involve expiration and renewal.

Method 1: Renew the subscription in the ApsaraDB for POLARDB console

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region in the upper-left corner to view all the clusters that you deploy in this region.

3. Find the target cluster, click the More icon in the Actions column, and choose Renew from the shortcut menu.

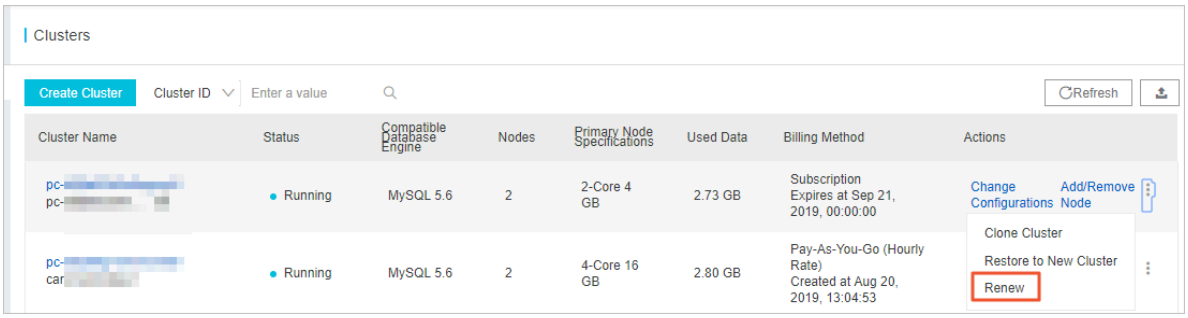

4. Specify the renewal duration, select the service agreement, and click Pay.

Method 2: Renew the subscription in the Renew console

1. Log on to the ApsaraDB for POLARDB console.
2. In the upper-right corner of the console, choose Billing Management > Renew.

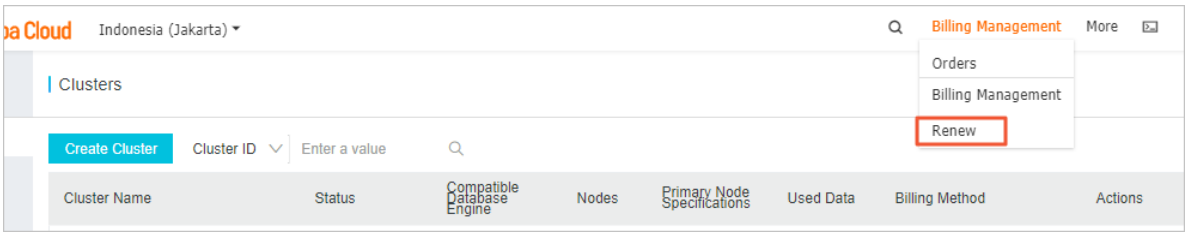

3. In the left-side navigation pane, click ApsaraDB for POLARDB.
4. Click the Manually Renew tab. Set the filtering conditions to find the target cluster. Click Renew in the Actions column corresponding to the cluster.

> **Note:**
> To enable manual renewal for a cluster on the Auto-Renew or Don't Renew tab, click Enable Manual Renew, and then click OK in the dialog box that appears.

5. Specify the renewal duration, select the service agreement, and click Pay.

Enable automatic renewal

If you enable automatic renewal, you will be free from regular manual renewal operations and concerns of service interruptions. For more information, see #unique_20.

5.3 Automatically renew the subscription to a cluster

A subscription-based cluster has a validity period. If the cluster is not renewed in a timely manner, service interruptions or even data loss will occur after it expires

. If you enable automatic renewal, you will be free from regular manual renewal operations and concerns of service interruptions.

> Note:
> Clusters purchased through the pay-as-you-go (hourly rate) billing method do not involve expiration and renewal.

Precautions

- Automatic fee deduction will begin nine days prior to the expiration of the cluster, supporting cash and coupons. Keep your account balance adequate.
- If you manually renew the cluster before the automatic deduction, the system will automatically renew the cluster nine days prior to the next expiration.
- The automatic renewal feature takes effect the next day after it is enabled. If your cluster expires the next day, renew it manually to prevent service interruptions. For more information, see #unique_22.

Enable automatic renewal when purchasing a cluster

> Note:
> After you enable automatic renewal, the system will automatically renew the subscription based on the subscription period. For example, if you purchase a cluster for three months and select automatic renewal, you will be charged a fee of the three-month subscription upon each automatic renewal.

When creating a cluster, you can select Auto Renew.

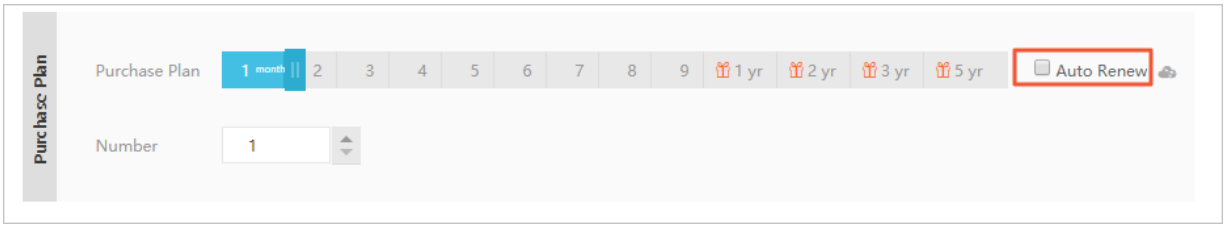

Enable automatic renewal after purchasing a cluster

> Note:
> After you enable automatic renewal, the system will automatically renew the subscription based on the renewal cycle you select. For example, if you select a three-month renewal cycle, you will be charged a fee of the three-month subscription upon each automatic renewal.

1. Log on to the ApsaraDB for POLARDB console.
2. In the upper-right corner of the console, choose Billing Management > Renew.

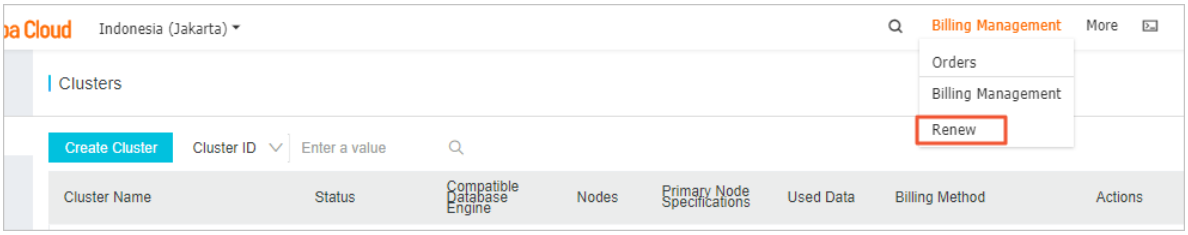

3. In the left-side navigation pane, click ApsaraDB for POLARDB.
4. Click the Manually Renew or Don't Renew tab on the Renew console. Set the filtering conditions to find the target cluster. Click Enable Auto-Renew in the Actions column corresponding to the cluster.
5. In the dialog box that appears, select the automatic renewal cycle, and click Enable Auto-Renew.

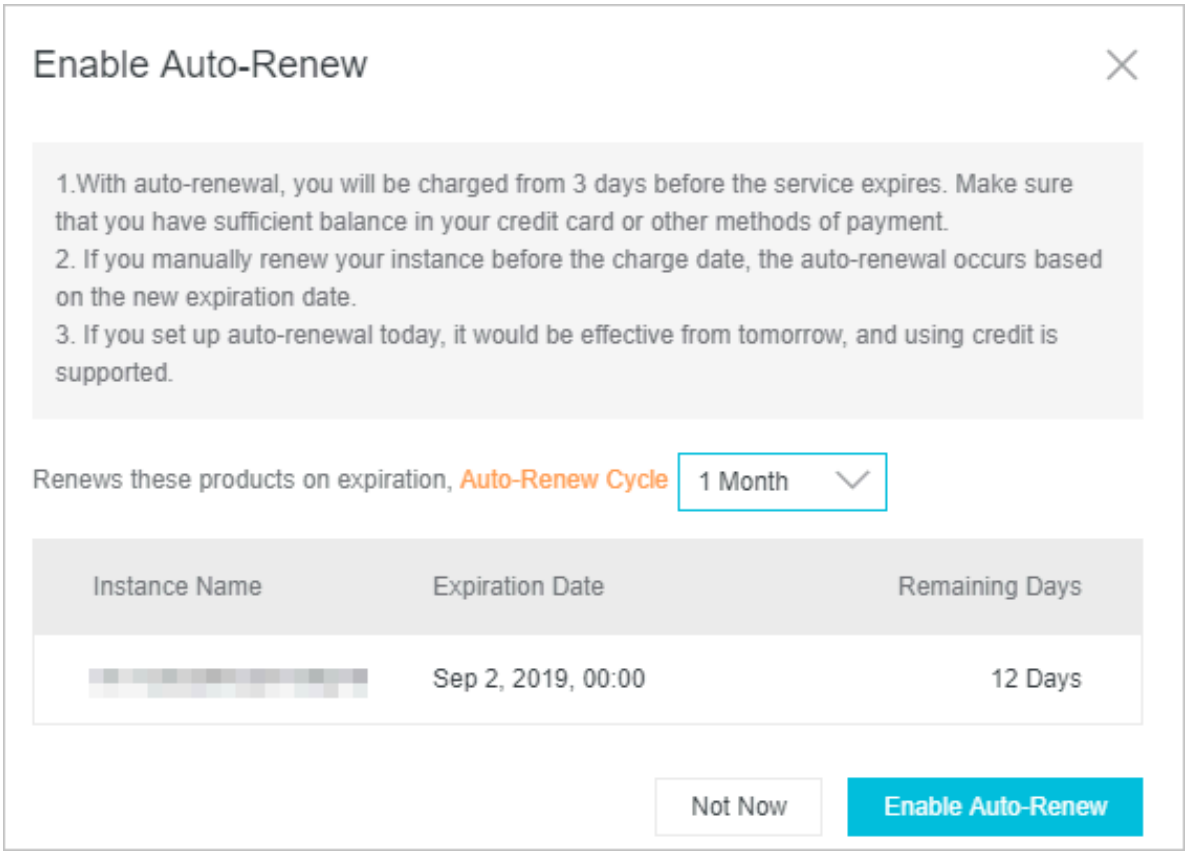

Edit the automatic renewal cycle

1. Log on to the ApsaraDB for POLARDB console.

2. In the upper-right corner of the console, choose Billing Management > Renew.

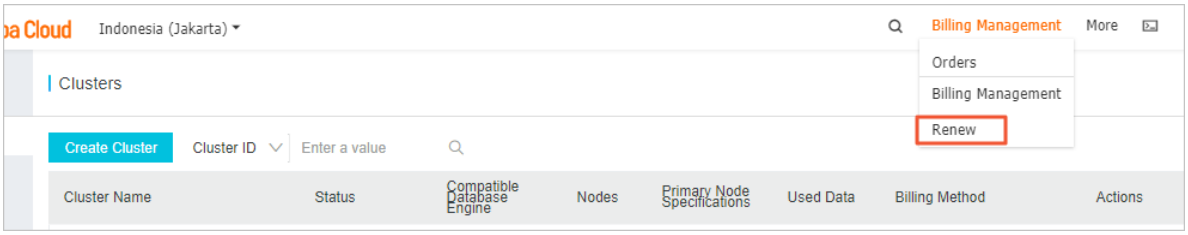

3. In the left-side navigation pane, click ApsaraDB for POLARDB.
4. Click the Auto-Renew tab on the Renew console. Set the filtering conditions to find the target cluster. Click Enable Auto-Renew in the Actions column corresponding to the cluster.
5. Click the Auto tab. Set the filtering conditions to find the target cluster. Click Modify Auto-Renew in the Actions column corresponding to the cluster.
6. In the dialog box that appears, edit the automatic renewal cycle, and click OK.

Disable automatic renewal

1. Log on to the ApsaraDB for POLARDB console.
2. In the upper-right corner of the console, choose Billing Management > Renew.

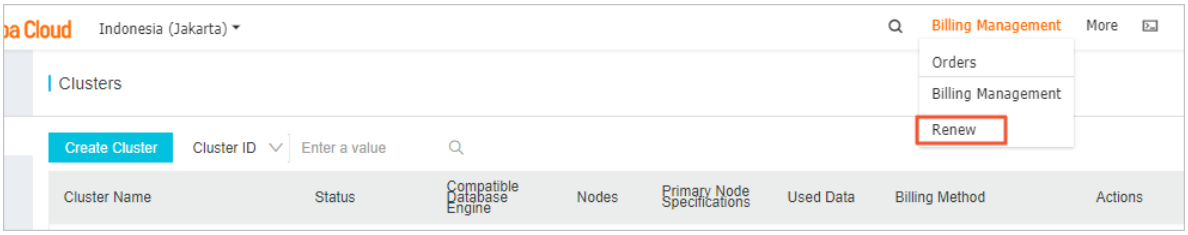

3. In the left-side navigation pane, click ApsaraDB for POLARDB.
4. Click the Auto-Renew tab on the Renew console. Set the filtering conditions to find the target cluster. Click Modify Auto-Renew in the Actions column corresponding to the cluster.

5. Select Disable Auto-Renew and click OK.

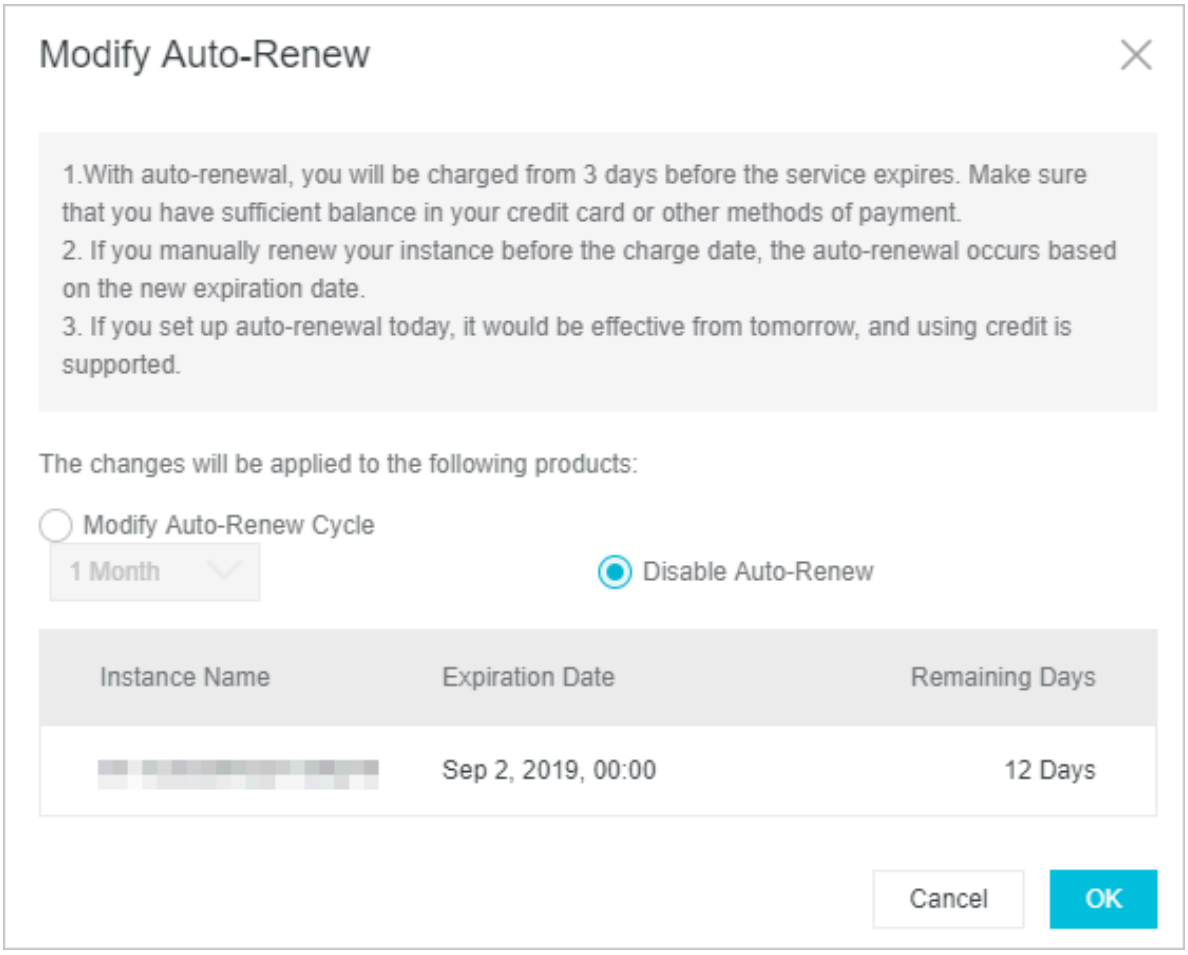

Related operations

Operation	Description
#unique_23	Creates a POLARDB cluster. Note: You can enable automatic renewal when you create a cluster.
#unique_24	Enables automatic renewal for a subscription-based cluster. Note: You can enable automatic renewal after you create a cluster.
#unique_25	Queries the automatic renewal status of a subscription-based cluster.

6 Connect to POLARDB

6.1 View connection endpoints

A POLARDB cluster includes cluster connection endpoints and primary connection endpoints.

Procedure

1. Log on to the POLARDB console.
2. Find the target POLARDB cluster and click its ID.
3. In the Access Information on the Basics page, view the connection endpoints of the POLARDB cluster.

Cluster and primary connection endpoints

Type	Description	Supported network type
Cluster connection endpoint (recommended)	An application only needs to connect to a cluster connection endpoint, then it can connect to all the nodes in the POLARDB cluster. The cluster endpoint supports read/write splitting. It sends write requests to the primary node and read requests to the primary and read-only nodes and can automatically balance load among these nodes. Note: The POLARDB cluster contains one default cluster connection endpoint. You can customize one or more cluster connection endpoints as needed. A custom connection endpoint can connect to specified nodes and work in the specified read/write mode. For more information, see #unique_28.	VPC and public network

Type	Description	Supported network type
Primary connection endpoint	A primary connection endpoint always connects to the primary node and supports read and write operations. If the primary node becomes faulty, the primary connection endpoint is automatically switched to the read-only node that is promoted to the primary node.	VPC and public network

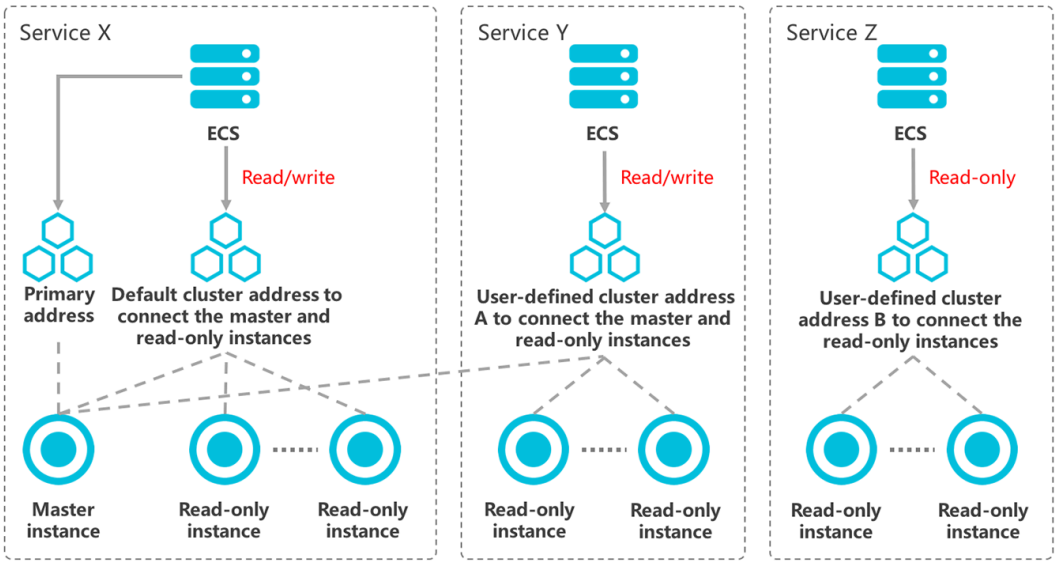

Private and public connection endpoints

Type	Description	Application scenario
Private connection endpoint	· Accessing the POLARDB cluster through a private connection endpoint maximizes performance. · A private connection endpoint cannot be released.	For example: · If your ECS instance is located in the same VPC as the POLARDB cluster , then your ECS instance can communicate with the POLARDB cluster through the VPC. · You can access the POLARDB cluster through a VPC by using DMS.

Type	Description	Application scenario
Public connection endpoint	· You must manually apply for a public connection endpoint, which can be released. · The public network refers to the Internet. Accessing the POLARDB cluster through the Internet cannot maximize performance.	For example, you can access the POLARDB cluster through the Internet to perform maintenance.

APIs

API	Description
#unique_7	Used to list the connection endpoints of a POLARDB cluster.
#unique_5	Used to create a public connection endpoint for a POLARDB cluster.
#unique_9	Used to change the default connection endpoint of a POLARDB cluster.
#unique_10	Used to release a cluster connection endpoint for a POLARDB cluster.

6.2 Set or release a custom cluster connection endpoint

The system automatically generates a default connection endpoint for a cluster. You can also manually add custom cluster connection endpoints.

You can use custom connection endpoints on a POLARDB for MySQL cluster. You can also set their read/write mode and consistency level and select associated read-only nodes based on different business scenarios. This enhances business flexibility.

Note:

· A cluster can have a maximum of four cluster connection endpoints, including one default connection endpoint and three custom connection endpoints.
· The default cluster connection endpoint cannot be released. A custom cluster connection endpoint can be released.

- Like a custom cluster connection endpoint, the default cluster connection endpoint supports custom settings. For more information, see Modify a custom cluster connection endpoint.

Prerequisites

You can directly add custom connection endpoints for clusters created on and after April 29, 2019. For clusters created before April 29, 2019, you must open a ticket before you can add custom cluster connection endpoints.

Add a custom cluster connection endpoint

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the Access Information section on the Basics page, click Create Custom Connection Endpoint next to Cluster Connection Endpoints (Recommended).

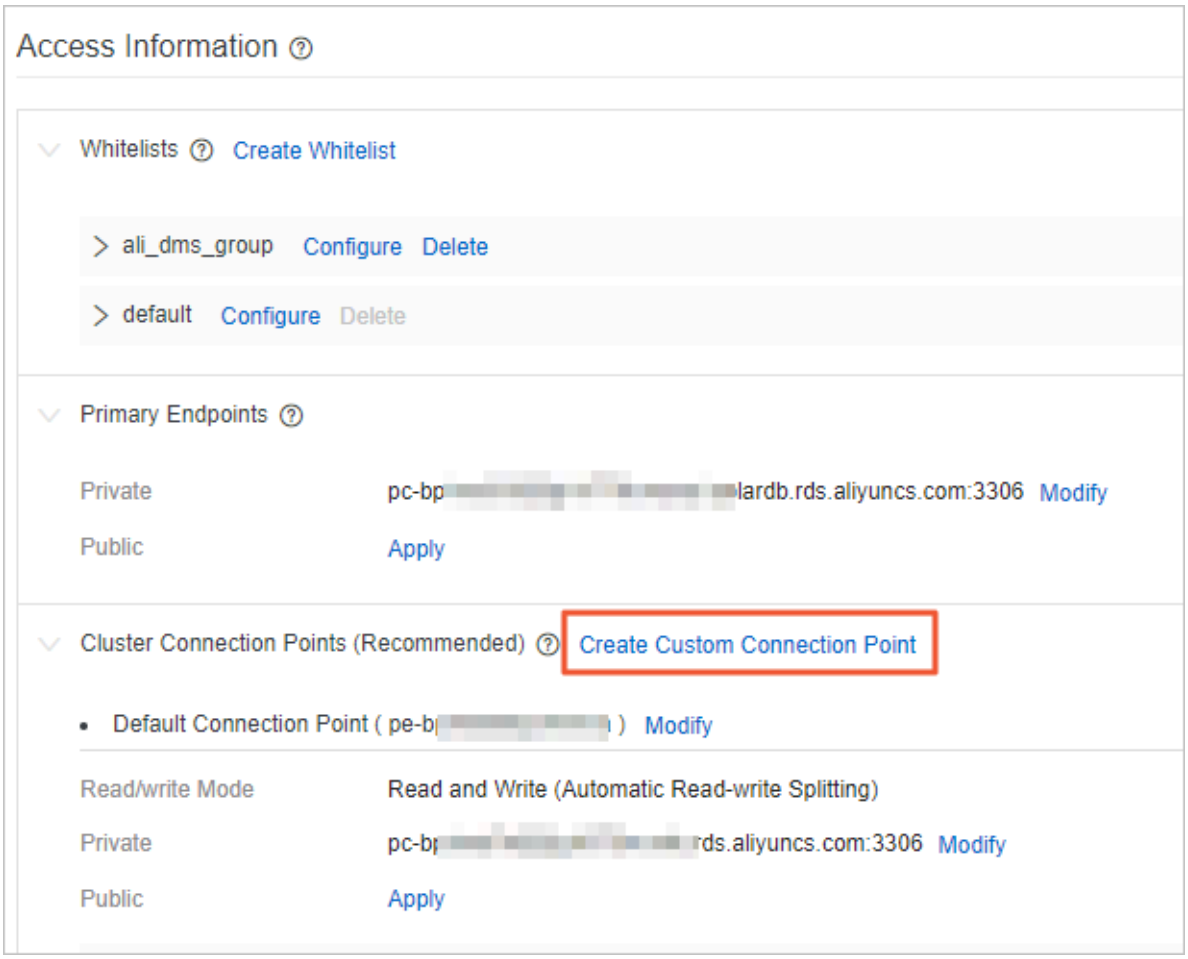

5. In the dialog box that appears, set parameters for creating a custom cluster connection endpoint. The following table describes the parameters.

Parameter	Description
Read/write Mode	Select the read/write mode of the connection endpoint. Valid values: Read Only and Read and Write (Automatic Read-write Splitting). Note: You can also modify the read/write mode of a custom connection endpoint after it is created. The modification takes effect only for the newly created connections. The existing connections remain in the original read/write mode.
Reader Nodes	From the Unselected Nodes list on the left, select the nodes that you want to add to the connection endpoint to process read requests. The available nodes include the primary node and all read-only nodes. The connection endpoint only sends read requests to the selected nodes. Note: · Select at least two nodes. · Write requests are sent only to the primary node regardless of whether the primary node is selected.
Automatically Add New Nodes	Specify whether a newly added node will be automatically added to the connection endpoint.
Load Balancing Policy	The scheduling policy for read requests among multiple read-only nodes when read-write splitting is enabled. The value is fixed.

Parameter	Description
Consistency Level	· Eventual Consistency: provides the best performance. · Session Consistency: guarantees the read consistency at the session level. In this mode, the load of the primary node is slightly increased. Note: If the read/write mode is set to Read Only, the value is fixed to Eventual Consistency.

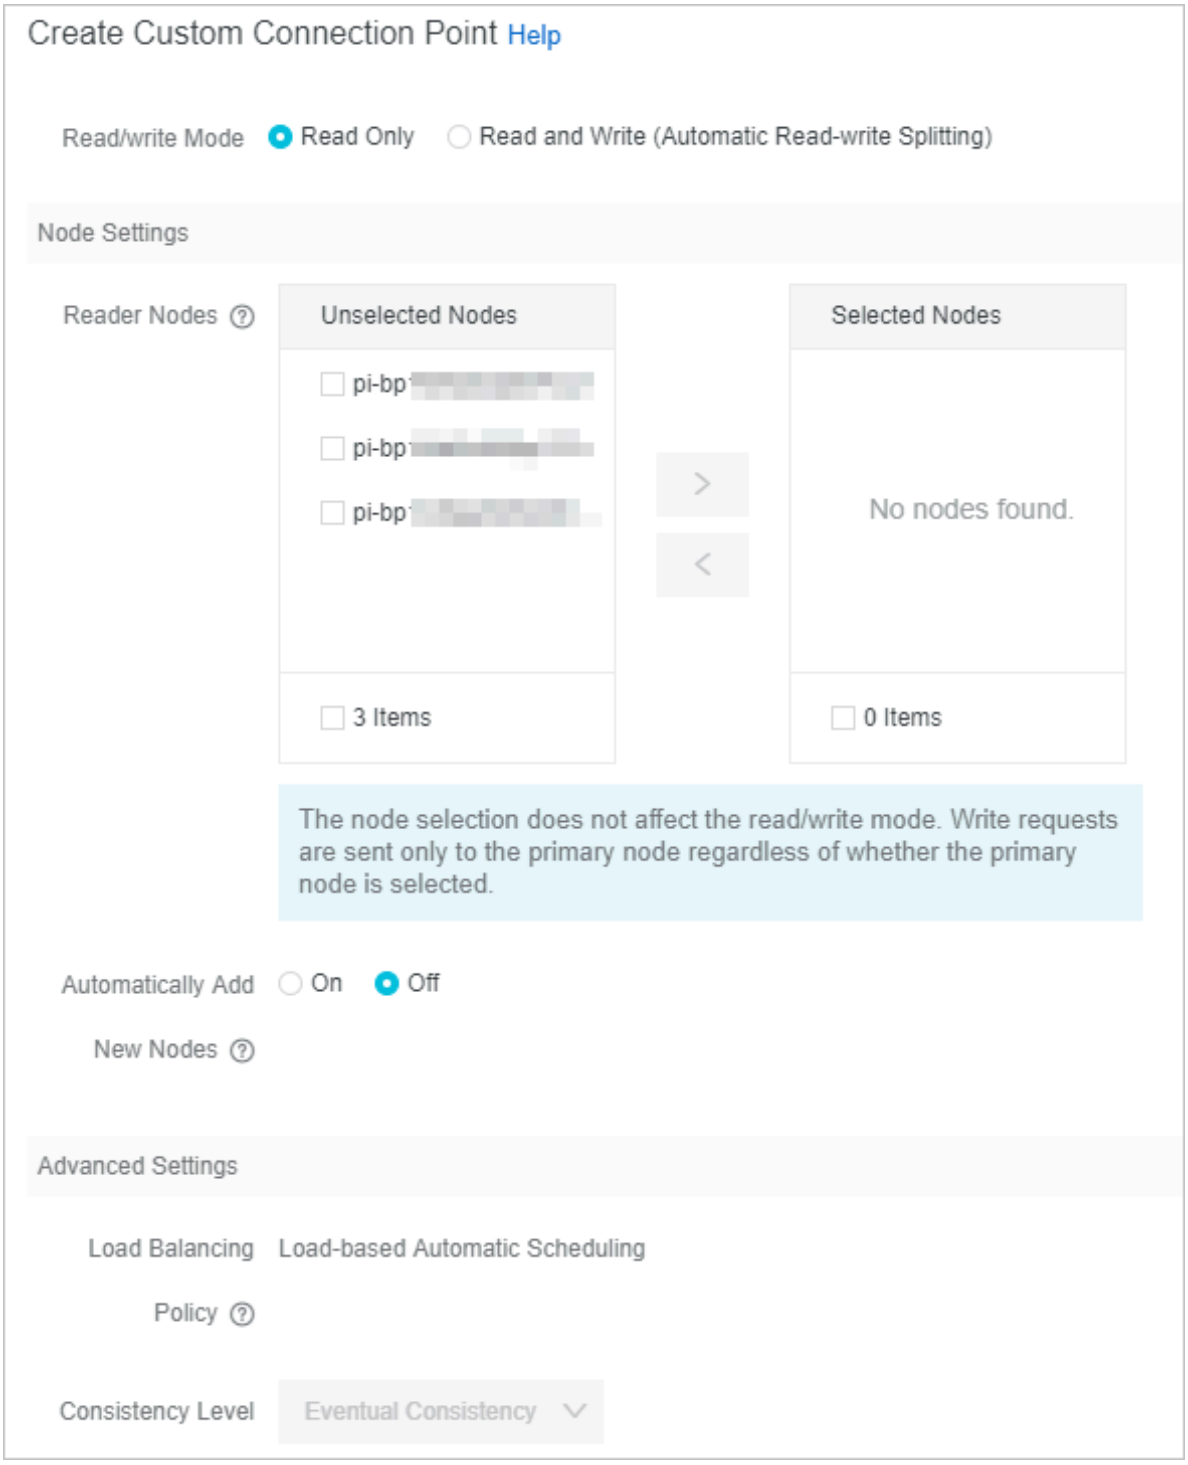

6. Click OK.

Modify a custom cluster connection endpoint

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.

4. In the Access Information section on the Basics page, click Modify next to a custom connection endpoint.

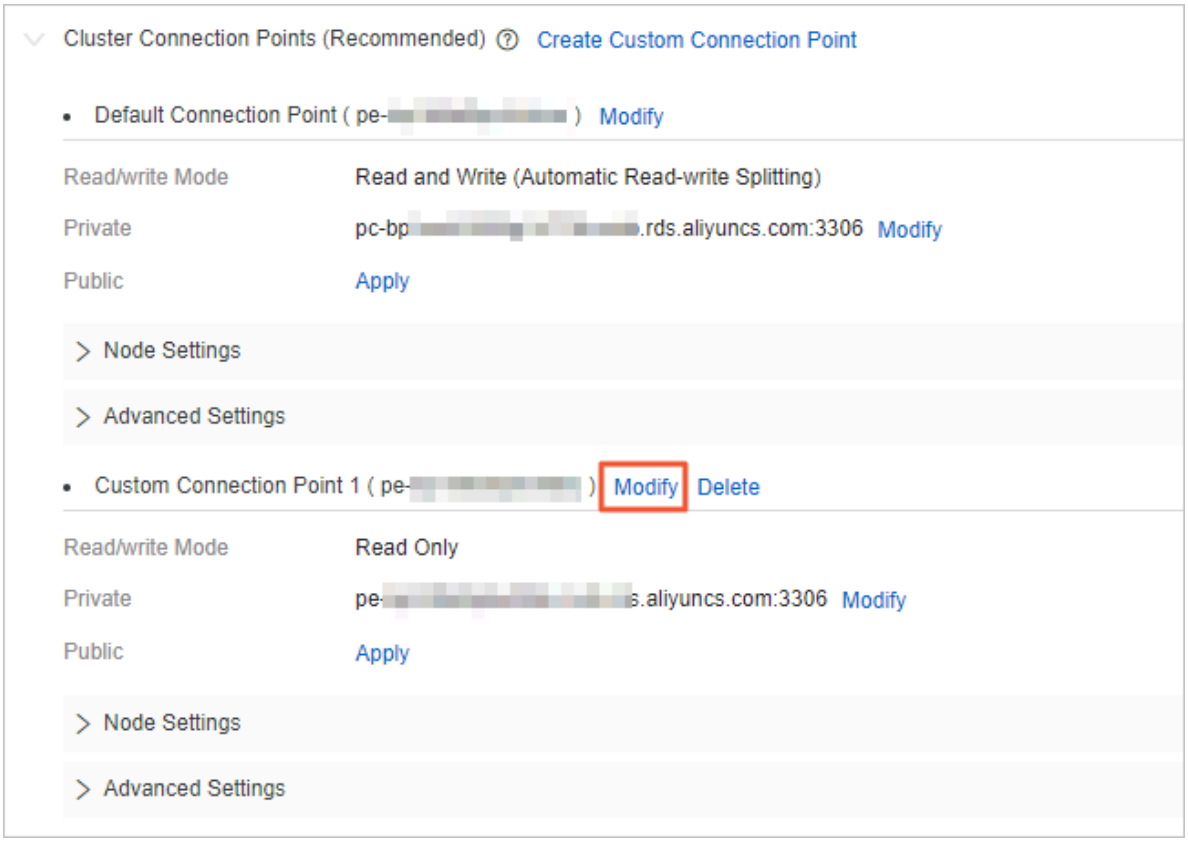

5. In the dialog box that appears, set parameters for modifying a custom cluster connection endpoint. The following table describes the parameters.

Parameter	Description
Read/write Mode	Select the read/write mode of the connection endpoint. Valid values: Read Only and Read and Write (Automatic Read-write Splitting).

Parameter	Description
Reader Nodes	From the Unselected Nodes list on the left, select the nodes that you want to add to the connection endpoint to process read requests. The available nodes include the primary node and all read-only nodes. The connection endpoint only sends read requests to the selected nodes. Note: · Select at least two nodes. · Write requests are sent only to the primary node regardless of whether the primary node is selected. · Adding nodes to a connection endpoint does not affect the use of the connection endpoint. However, when a node is removed from the connection endpoint, the persistent connection on the node is interrupted.
Automatically Add New Nodes	Specify whether a newly added node will be automatically added to the connection endpoint.
Load Balancing Policy	The scheduling policy for read requests among multiple read-only nodes when read-write splitting is enabled. The value is fixed.
Consistency Level	· Eventual Consistency: provides the best performance. · Session Consistency: guarantees the read consistency at the session level. In this mode, the load of the primary node is slightly increased. Note: · If the read/write mode is set to Read Only, the value is fixed to Eventual Consistency. · The modification of the consistency level immediately takes effect for all connections.

6. Click OK.

Release a custom cluster connection endpoint

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.

4. In the Access Information section on the Basics page, find the target custom connection endpoint under Cluster Connection Endpoints (Recommended), and click Delete.

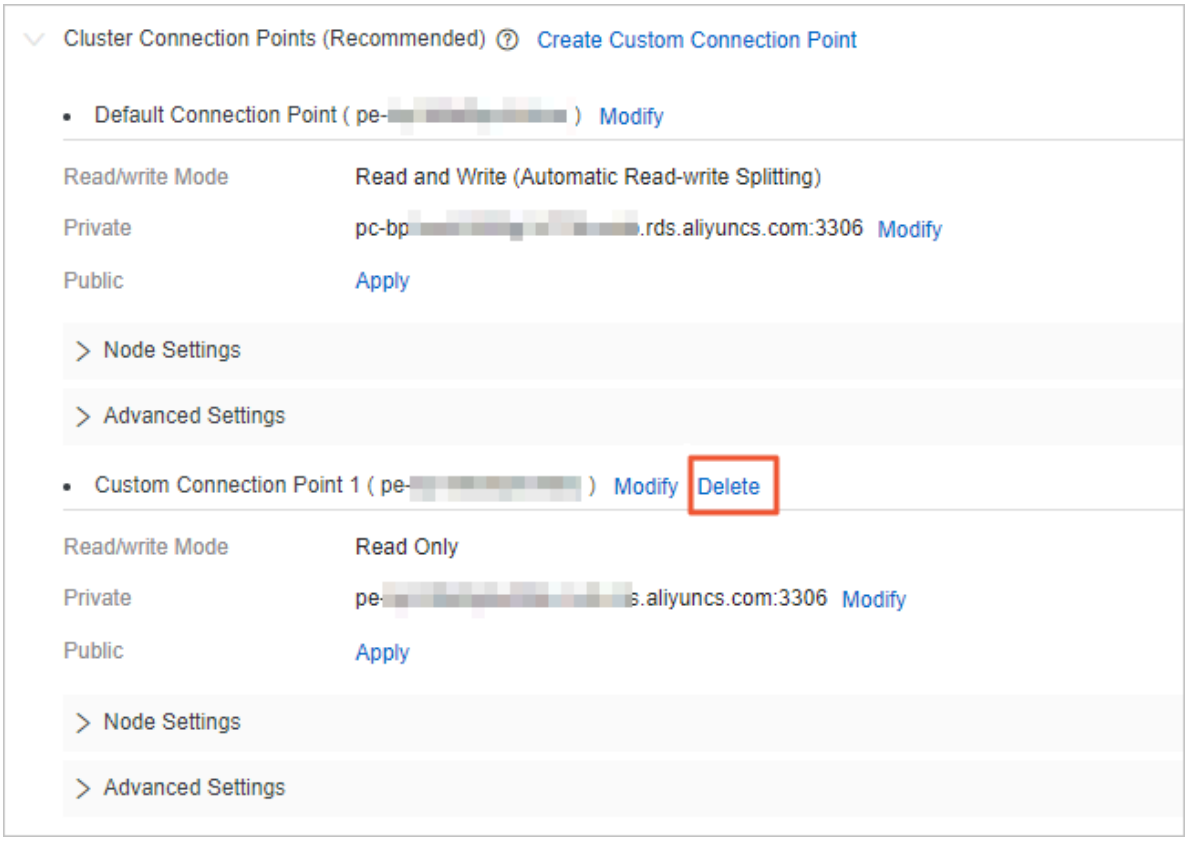

5. In the dialog box that appears, click OK.

Related API operations

API operation	Description
#unique_6	Creates a custom cluster connection endpoint.
#unique_7	Queries cluster connection endpoints.
#unique_8	Modifies a cluster connection endpoint.
#unique_11	Releases a custom cluster connection endpoint.

6.3 Connect to a POLARDB for MySQL cluster

This topic describes how to connect to a POLARDB for MySQL cluster through Data Management Service (DMS), through a general-purpose database client, and through the CLI.

Prerequisites

A privileged or standard account has been created for the POLARDB cluster. For more information, see #unique_31.

Connect to a POLARDB cluster through DMS

DMS is a graphical data management tool provided by Alibaba Cloud. It offers an integrated solution for data management, structure management, access security , BI charts, data trends, data tracking, performance management, performance management. With DMS, you can manage relational databases (including MySQL, SQL Server, and PostgreSQL), NoSQL databases (including MongoDB and Redis), and Linux servers.

1. Find the target POLARDB cluster and click its ID.
2. In the upper-right corner, click Log On to Database.

3. On the displayed logon page, enter the primary connection endpoint and port number, which are separated by a comma (,), enter the username and password of the privileged or standard account, and click Log On.

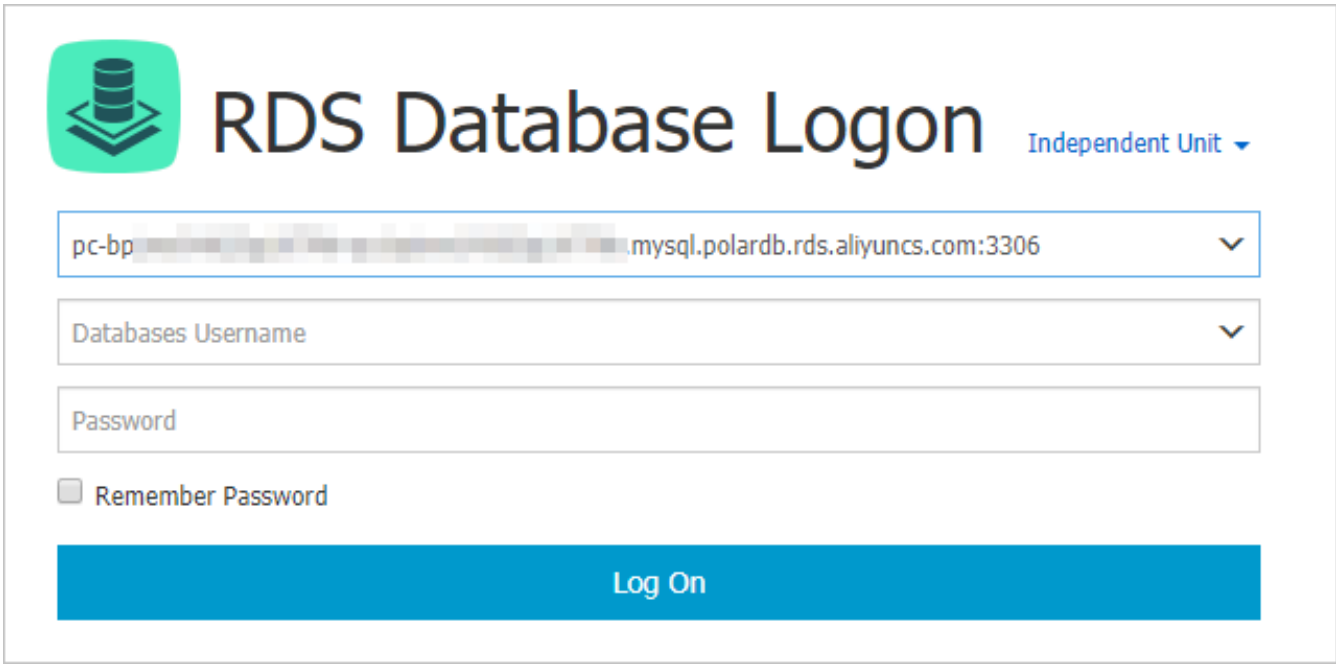

Note:
You can log on to your POLARDB cluster through DMS only by using the primary connection endpoint but not a cluster connection endpoint. For more information about how to view connection endpoints, see #unique_32.

Connect to a POLARDB cluster through a database client

You can use any general-purpose database client to connect to a POLARDB cluster. This topics uses the HeidiSQL database client as an example.

1. Start HeidiSQL.

2. In the lower-left corner, click New.

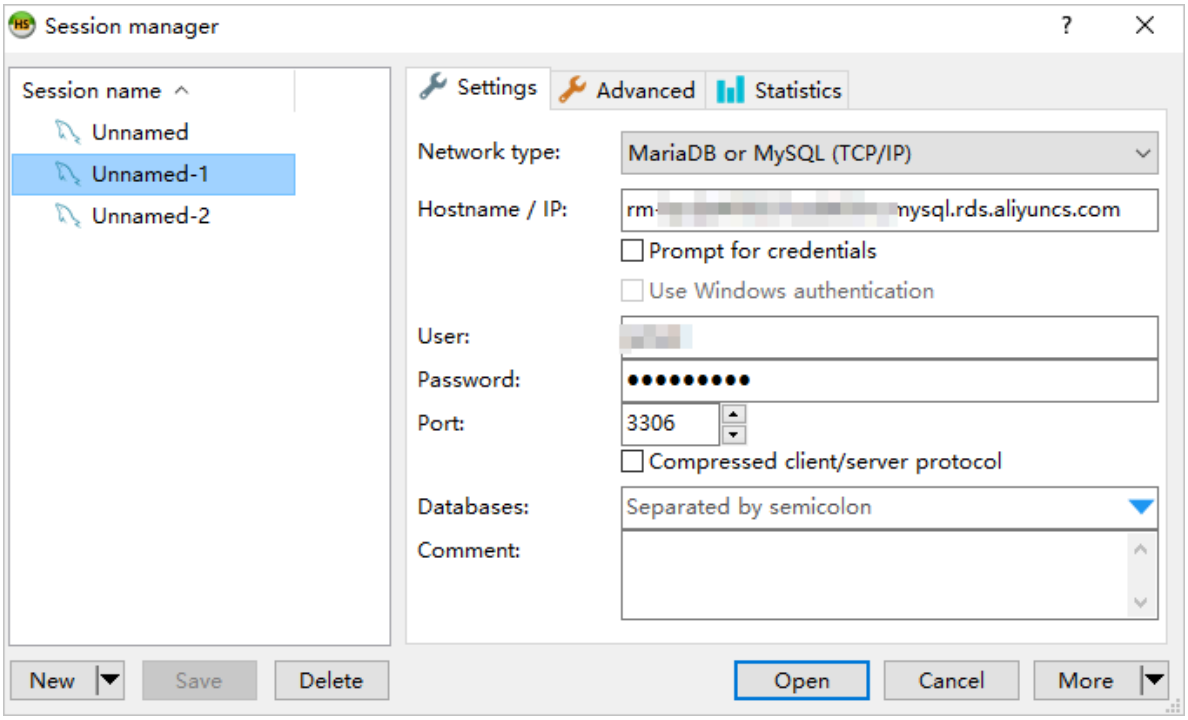

3. Enter the POLARDB cluster information.

Parameter	Description
Network type	The method of connecting to the POLARDB cluster. Select MariaDB or MySQL (TCP/IP).

Parameter	Descripition
Hostname/IP	The private or public connection endpoint of the POLARDB cluster. · If the database client is deployed in an ECS instance and the region and network type of the ECS instance are the same as those of the POLARDB cluster, use the private connection endpoint to establish a secure, efficient connection. · In the other situations, use the public connection endpoint. To view the private and public connection endpoints of the POLARDB cluster, follow these steps: a. Log on to the POLARDB console. b. In the upper-left corner, select the region where the target POLARDB cluster is located. c. Find the target POLARDB cluster and click its ID. d. On the Basics page, view the private and public connection endpoints and ports.
User	The username of the account you use to access the POLARDB cluster.
Password	The password of the account you use to access the POLARDB cluster.
Port	The port number corresponding to the private or public connection endpoint of the POLARDB cluster.

4. Click Open. If the connection information is correct, the POLARDB cluster gets connected.

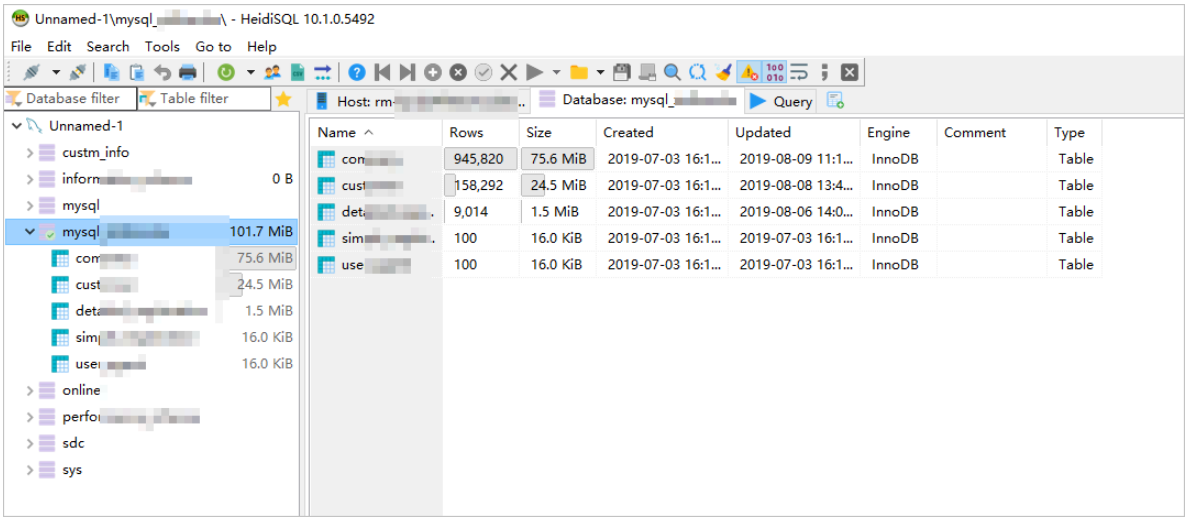

Connect to a POLARDB cluster through the CLI

If MySQL is installed on your server, you can run the following command to connect to a POLARDB for MySQL cluster:

```
mysql -h <Connection endpoint> -P <Port number> -u <Username> -p <Password> -D <Name of the POLARDB cluster>
```

Parameter	Description	Example
-h	The private or public connection endpoint of the POLARDB cluster. For more information, see #unique_33.	`pc-bpxxxxxxxx xxxxxx.mysql.polardb.rds.aliyuncs.com`

Parameter	Description	Example
-P	The port number of the POLARDB cluster. · If you use a private connection endpoint, enter the private port number. · If you use a public connection endpoint, enter the public port number Note: · The default port number is 3306. · If the default port number is used for connecting to the POLARDB cluster, you can leave this parameter blank.	`3306`
-u	The username of the account that you use to access the POLARDB cluster.	`root`
-p	The password of the account that you use to access the POLARDB cluster. Note: This parameter is optional. · If you do not set this parameter, the system will ask you to enter the password later. · If you set this parameter, do not leave any spaces between `-p` and the entered password.	`password23 3`
-D	The name of the POLARDB cluster to which you want to log on. Note: · This parameter is optional. · You can remove `-D` and enter only the cluster name.	`mysql`

```
[root@iZbp1454s3dm0vndsljexlZ ~]# mysql -hpc-[redacted].mysql.polardb.rds
.aliyuncs.com -u[redacted] -p mysql
Enter password:
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Welcome to the MariaDB monitor.  Commands end with ; or \g.
Your MySQL connection id is 1409
Server version: 5.6.16-log Source distribution

Copyright (c) 2000, 2018, Oracle, MariaDB Corporation Ab and others.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

MySQL [mysql]>
```

6.4 Connect to a POLARDB for MySQL cluster

This topic describes how to connect to a POLARDB for MySQL cluster through Data Management Service (DMS), through a general-purpose database client, and through the CLI.

Prerequisites

A privileged or standard account has been created for the POLARDB cluster. For more information, see #unique_31.

Connect to a POLARDB cluster through DMS

DMS is a graphical data management tool provided by Alibaba Cloud. It offers an integrated solution for data management, structure management, access security , BI charts, data trends, data tracking, performance management, performance management. With DMS, you can manage relational databases (including MySQL, SQL Server, and PostgreSQL), NoSQL databases (including MongoDB and Redis), and Linux servers.

1. Find the target POLARDB cluster and click its ID.
2. In the upper-right corner, click Log On to Database.

3. On the displayed logon page, enter the primary connection endpoint and port number, which are separated by a comma (,), enter the username and password of the privileged or standard account, and click Log On.

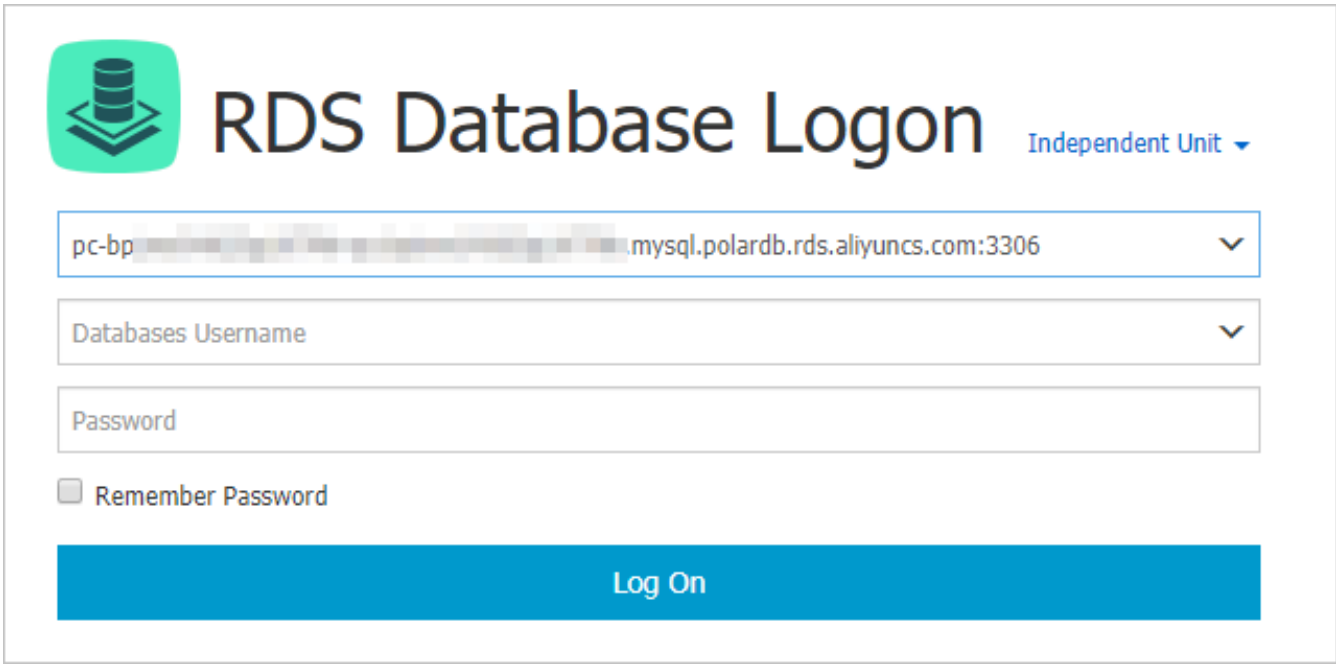

> **Note:**
> You can log on to your POLARDB cluster through DMS only by using the primary connection endpoint but not a cluster connection endpoint. For more information about how to view connection endpoints, see #unique_32.

Connect to a POLARDB cluster through a database client

You can use any general-purpose database client to connect to a POLARDB cluster. This topics uses the HeidiSQL database client as an example.

1. Start HeidiSQL.

2. In the lower-left corner, click New.

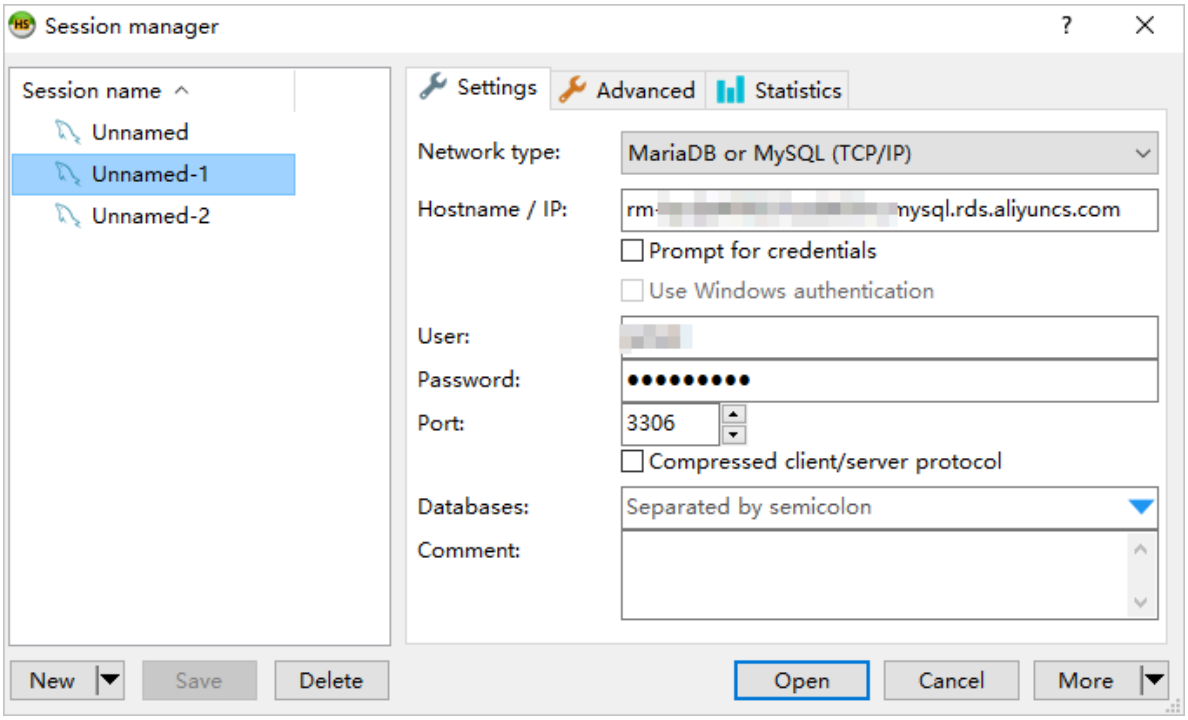

3. Enter the POLARDB cluster information.

Parameter	Description
Network type	The method of connecting to the POLARDB cluster. Select MariaDB or MySQL (TCP/IP).

Parameter	Descripition
Hostname/IP	The private or public connection endpoint of the POLARDB cluster. · If the database client is deployed in an ECS instance and the region and network type of the ECS instance are the same as those of the POLARDB cluster, use the private connection endpoint to establish a secure, efficient connection. · In the other situations, use the public connection endpoint. To view the private and public connection endpoints of the POLARDB cluster, follow these steps: a. Log on to the POLARDB console. b. In the upper-left corner, select the region where the target POLARDB cluster is located. c. Find the target POLARDB cluster and click its ID. d. On the Basics page, view the private and public connection endpoints and ports.
User	The username of the account you use to access the POLARDB cluster.
Password	The password of the account you use to access the POLARDB cluster.
Port	The port number corresponding to the private or public connection endpoint of the POLARDB cluster.

4. Click Open. If the connection information is correct, the POLARDB cluster gets connected.

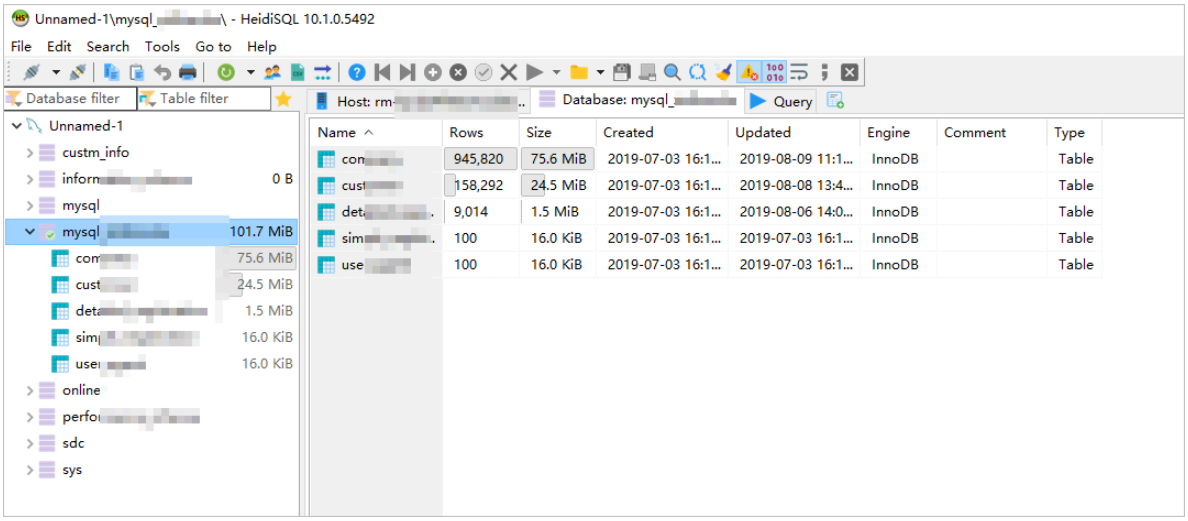

Connect to a POLARDB cluster through the CLI

If MySQL is installed on your server, you can run the following command to connect to a POLARDB for MySQL cluster:

```
mysql -h <Connection endpoint> -P <Port number> -u <Username> -p <Password> -D <Name of the POLARDB cluster>
```

Parameter	Description	Example
-h	The private or public connection endpoint of the POLARDB cluster. For more information, see #unique_33.	pc-bpxxxxxxxx xxxxxx.mysql.polardb.rds.aliyuncs.com

Parameter	Description	Example
-P	The port number of the POLARDB cluster. · If you use a private connection endpoint, enter the private port number. · If you use a public connection endpoint, enter the public port number Note: · The default port number is 3306. · If the default port number is used for connecting to the POLARDB cluster, you can leave this parameter blank.	`3306`
-u	The username of the account that you use to access the POLARDB cluster.	`root`
-p	The password of the account that you use to access the POLARDB cluster. Note: This parameter is optional. · If you do not set this parameter, the system will ask you to enter the password later. · If you set this parameter, do not leave any spaces between `-p` and the entered password.	`password23 3`
-D	The name of the POLARDB cluster to which you want to log on. Note: · This parameter is optional. · You can remove `-D` and enter only the cluster name.	`mysql`

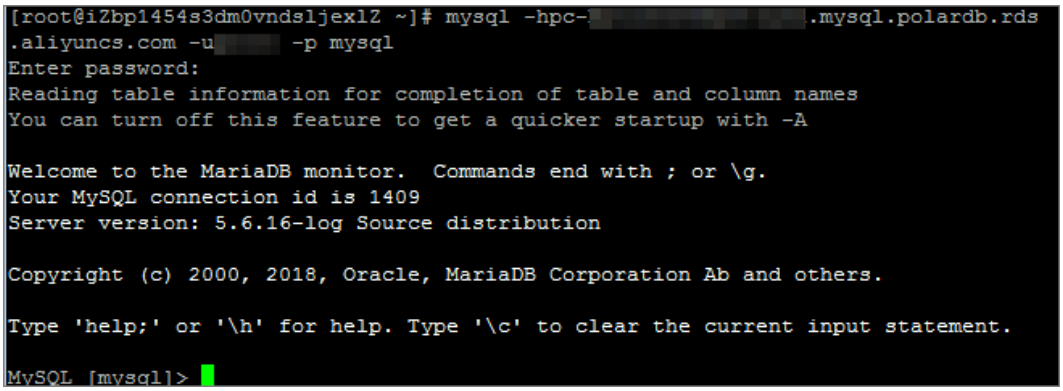
[root@iZbp1454s3dm0vndsljexlZ ~]# mysql -hpc-[redacted].mysql.polardb.rds
.aliyuncs.com -u[redacted] -p mysql
Enter password:
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A
Welcome to the MariaDB monitor. Commands end with ; or \g.
Your MySQL connection id is 1409
Server version: 5.6.16-log Source distribution
Copyright (c) 2000, 2018, Oracle, MariaDB Corporation Ab and others.
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
MySQL [mysql]>

7 Cluster management

7.1 Create a POLARDB for MySQL cluster

This topic describes how to create a POLARDB for MySQL cluster in the POLARDB console.

Prerequisites

You have registered an Alibaba Cloud account or created a Resource Access Management (RAM) user.

- Click here to register an Alibaba Cloud account.
- For more information about how to create and grant permissions to a RAM user, see #unique_36.

Context

A POLARDB for MySQL cluster contains one primary node and up to 15 read-only nodes. (At least one read-only node is required to provide active-active high availability support). A node is a virtual DB server, where you can create and manage one or more databases.

Note:

- POLARDB supports Virtual Private Clouds (VPCs) only. A VPC is an isolated network in Alibaba Cloud that is more secure than a classic network.
- We recommend that you use POLARDB with Elastic Compute Service (ECS) and place them in the same VPC to achieve optimal performance. If your ECS instance is created in a classic network, you must migrate it to a VPC.

Procedure

1. Log on to Alibaba Cloud.
   - Click here to log on with your Alibaba Cloud account.
   - Click here to log on with your RAM user. For more information, see #unique_36/unique_36_Connect_42_section_zb2_54q_tdb.
2. Log on to the POLARDB console.
3. Click Create Cluster.

4. Select Subscription or Pay-As-You-Go.
   - Subscription: Pay for the DB servers of the primary node and of a read-only node by selecting either a monthly or annual subscription. Storage consumed by your POLARDB database is billed per GB/hour increments, and your payment is deducted from your account on an hourly basis. The Subscription billing method is more cost-effective for long term use. You can save more with longer subscription periods.
   - Pay-As-You-Go: DB servers are billed per hour and storage consumed by your POLARDB database is billed per GB/hour based on actual increments. Your payment is deducted from your account on an hourly basis. We recommend that you select the Pay-As-You-Go billing method for short term use.
5. Set the following parameters.

Console section	Parameter	Description
Basic	Region	Select the region in which your cluster resides. You cannot change the region once you confirm your order. Note: The POLARDB cluster must be located in the same region as the ECS instance to be connected. Otherwise, you can connect the POLARDB cluster to the ECS instance only through the Internet, which may degrade performance.
	Primary Availability Zone	The primary zone of the POLARDB cluster. · Zones are independent physical areas in a region. The zones in the same region are basically the same. · The POLARDB cluster and the ECS instance to be connected can be located in the same zone or in different zones. · You only need to select a primary zone. The system automatically assigns a secondary zone.
	Network Type	· You do not need to specify the network type. · POLARDB supports VPCs only. A VPC is an isolated virtual network and is more secure than a classic network.

Console section	Parameter	Description
	VPC Network VSwitch	Make sure that you place your POLARDB cluster and the ECS instance to be connected in the same VPC. Otherwise, they cannot intercommunicate through the intranet and achieve optimal performance. · Select the VPC if you have created a VPC that meets your network plan. For example, if you have created an ECS instance and the VPC where it resides meets your network plan, select this VPC. · Alternatively, use the default VPC and VSwitch. - Default VPC: ■ It is a unique VPC in your selected region. ■ The network mask for a default VPC has 16 bits , such as 172.31.0.0/16, providing up to 65,536 private IP addresses. ■ It is not counted against the total number of VPCs that you can create. - Default VSwitch: ■ It is a unique VSwitch in your selected region. ■ The network mask for a default VSwitch has 20 bits, such as 172.16.0.0/20, providing up to 4,096 private IP addresses. ■ It is not counted against the total number of Vswitches that you can create in a VPC. · If the default VPC and VSwitch cannot satisfy your requirements, you can create your own VPC and VSwitch.
Instance	Database Engine	· Fully compatible with MySQL 8.0 · Fully compatible with MySQL 5.6 · Fully compatible with PostgreSQL 11 · Highly compatible with Oracle
	Node Specification	Select a specification for your database node according to your needs. All POLARDB nodes own exclusive resources.

Console section	Parameter	Description
	Number Nodes	· You do not need to specify the number. By default, the system creates one read-only node with the same specifications as the primary node. · If a primary node fails, the system automatically promotes the read-only node as a primary node and generates a new read-only node. · For more information about read-only nodes, see Architecture.
	Storage Cost	You do not need to specify the storage cost. The system calculates the fees by hour based on the storage usage. For more information, see #unique_37. Note: You do not need to specify the storage capacity of the POLARDB cluster. The storage capacity automatically scales based on the data volume.

6. Specify the Purchase Plan (only applicable to Subscription clusters) and Number, and click Buy Now.

   Note:
   You can create up to 50 POLARDB clusters at a time when, for example, you want to roll out games in batches.

7. On the Order Confirmation page, confirm your order information, select ApsaraDB for POLARDB Pay-As-You-Go Agreement of Service, and then click Pay Now.

   The POLARDB cluster will be created within 10 minutes after you complete the payment. You can view the POLARDB cluster in the cluster list, and the primary and read-only nodes of the POLARDB cluster in the node list.

   Note:

   · If some of the nodes are in the Runningstate, the POLARDB cluster is unavailable and is still being created. The POLARDB cluster is only available when the cluster status is Running.
   · Make sure that you have selected the correct region, or you cannot view the POLARDB cluster or the nodes in it.

APIs

API	Description
#unique_23	Used to create a POLARDB cluster.
#unique_38	Used to list POLARDB clusters.
#unique_39	Used to view the attributes of a POLARDB cluster.
#unique_25	Used to query the automatic renewal status of a POLARDB cluster that uses the Subscription billing method.
#unique_24	Used to set the automatic renewal status of a POLARDB cluster that uses the Subscription billing method.

7.2 Deploy a multi-zone cluster

POLARDB for MySQL supports deploying a cluster across multiple zones. Compared with single-zone clusters, multi-zone clusters have better disaster recovery capabilities and can withstand breakdowns in a data center.

Multi-zone architecture

When a multi-zone cluster is deployed, data is distributed across multiple zones. Currently, compute nodes must be deployed in the primary zone. ApsaraDB for POLARDB reserves sufficient resources in a secondary zone to ensure a successful failover when the primary zone fails. The following figure shows the multi-zone architecture.

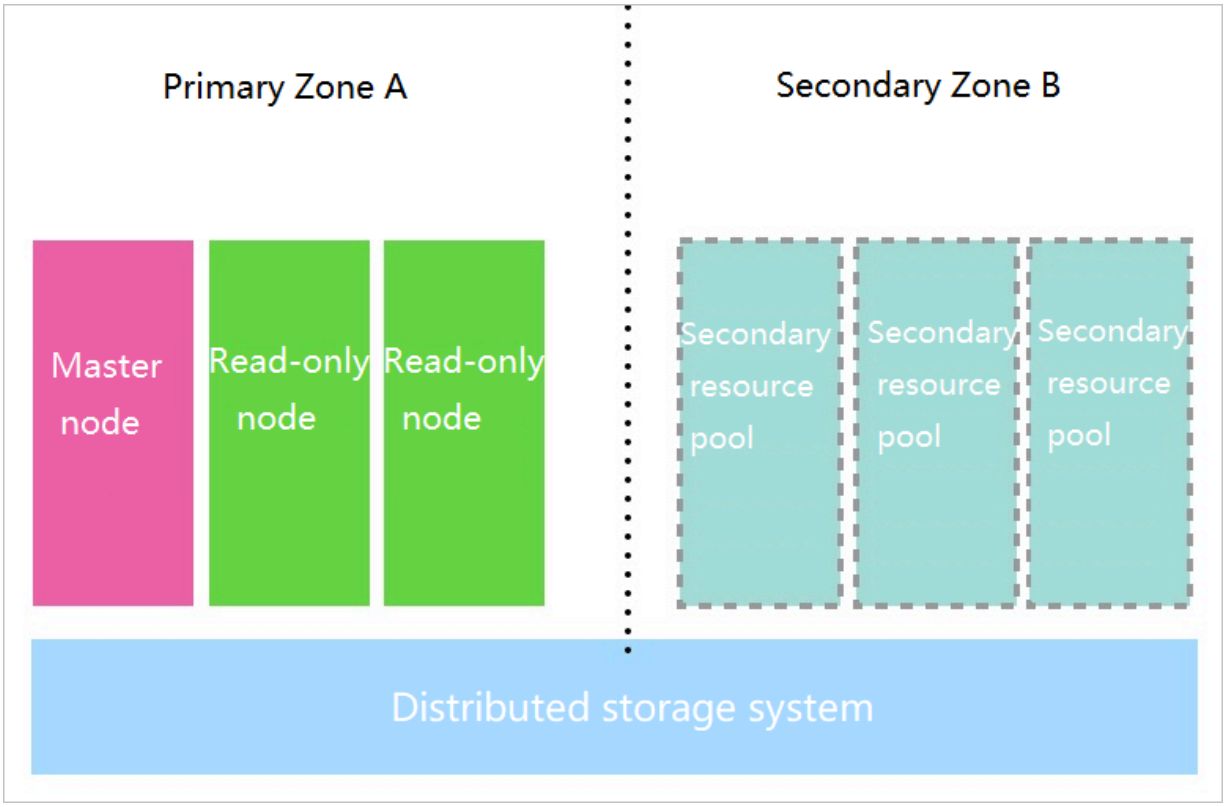

Pricing

No additional fee is required for multi-zone deployment.

> **Note:**
> You can also upgrade a single-zone cluster to a multi-zone cluster for free.

Prerequisites

- The cluster resides in China (Hangzhou) or China (Zhangjiakou).
- There are at least two zones with sufficient computing resources in the region.

Establish the multi-zone architecture

When the prerequisites are met, the cluster you create is set as a multi-zone cluster by default.

You can also upgrade the existing single-zone clusters to multi-zone ones. This upgrade is automatically completed through online data migration without affecting your business.

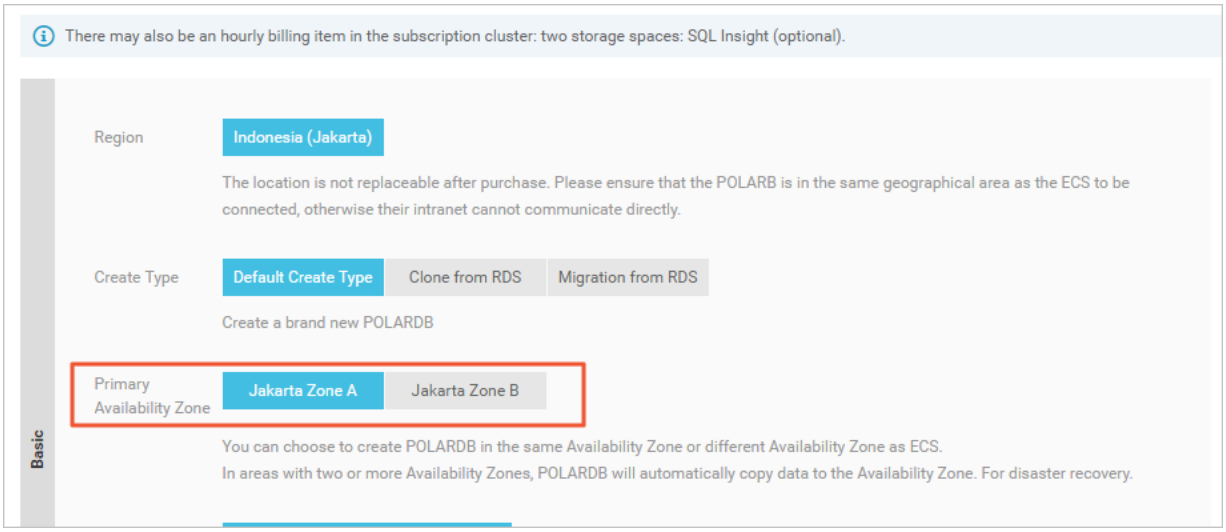

View the zones of a cluster

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the Basic Information section on the Basics page, view the zones of the cluster listed in Zones.

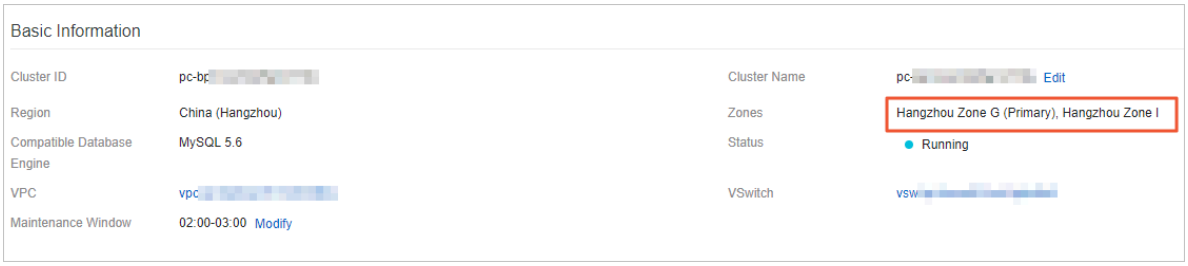

7.3 Add or remove a node

You can manually add or remove read-only nodes after creating an ApsaraDB for POLARDB cluster. An ApsaraDB for POLARDB cluster can contain a maximum of 15 read-only nodes. The cluster must have at least one read-only node to ensure high availability. All nodes in a cluster have the same specifications.

Impact of the node quantity on performance

For more information, see the POLARDB for MySQL performance white paper.

Node cost

The billing methods for adding nodes are as follows:

- If the cluster is charged in subscription mode, the added nodes are also charged in this mode.
- If the cluster is charged in pay-as-you-go mode (hourly rate), the added nodes are also charged in this mode.

> Note:
>
> - The read-only nodes that you purchase in either subscription or pay-as-you-go mode can be released at any time. After they are released, the system will refund or stop billing.
> - The added nodes are only charged based on the node specifications. For more information, see #unique_37. The storage fee is charged based on the actual data volume, regardless of the number of nodes.

Important notes

- You can add or remove read-only nodes only when the cluster does not have pending configuration change orders.
- To avoid misoperations, only one read-only node can be added or removed at a time. You need to repeat the same operations to add or remove multiple nodes.
- It takes about 5 minutes for the added or removed node to take effect.

Add a read-only node

> Note:
> After a read-only node is added, the newly created read-write splitting connection forwards requests to the node. The read-write splitting connection created before the new read-only node is added does not forward requests to the new read-only node. You need to disconnect and then re-establish the connection. For example, you can restart the application.

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.

3. Go to the Add/Remove Node page by using either of the following methods:
   - Find the target cluster and click Add/Remove Node in the Actions column.

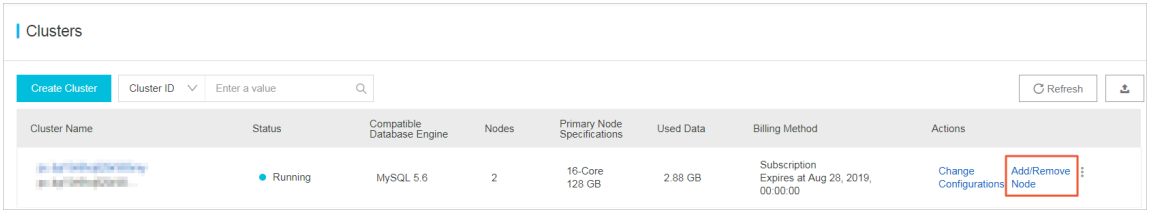

   - Find the target cluster, click the cluster ID, and then click Add/Remove Node in the Node Information section.

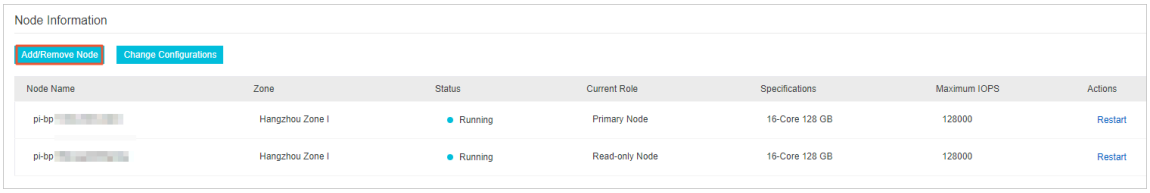

4. Select Add Node and click OK.

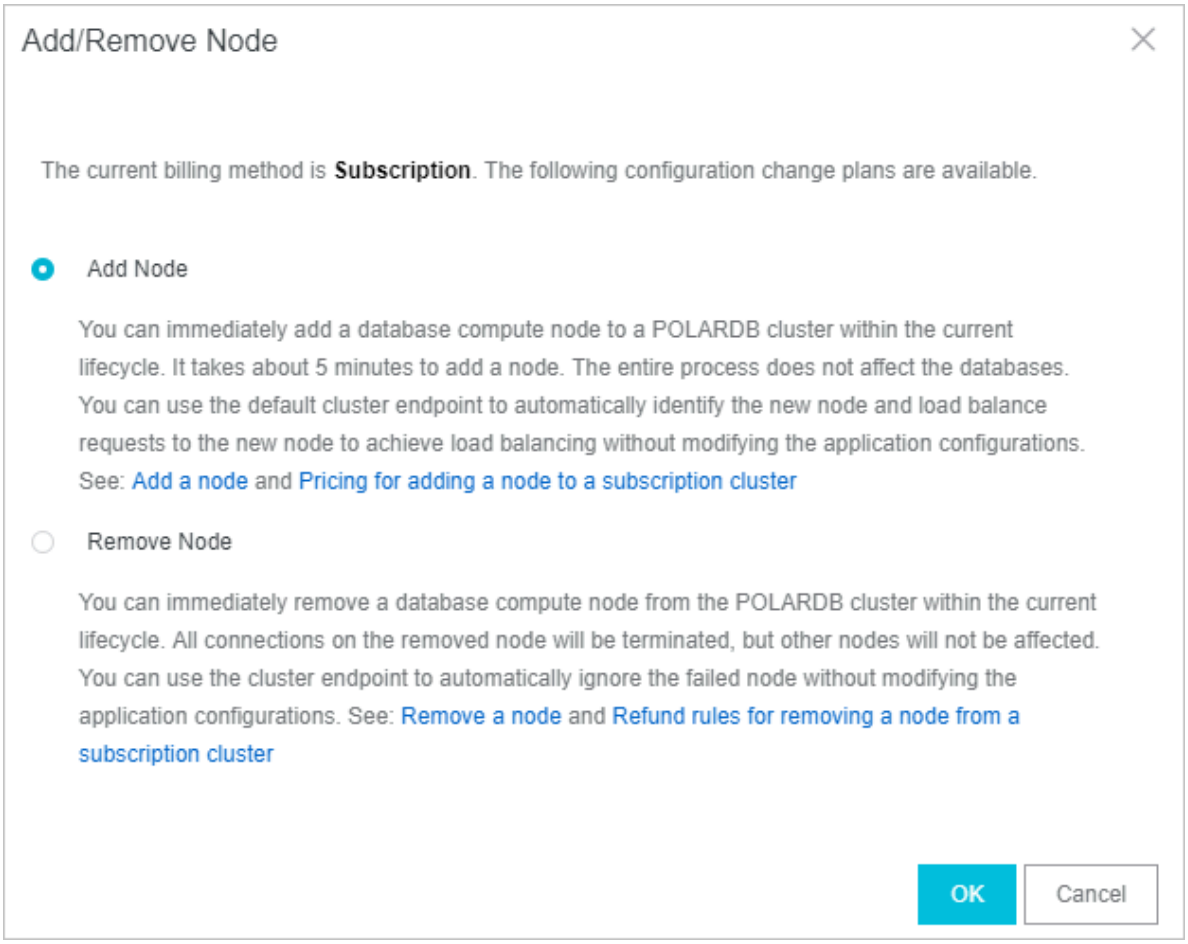

5. Click  to add a read-only node. Read and agree to the service agreement by selecting the check box, and click Pay to complete the payment.

Remove a read-only node

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Go to the Add/Remove Node page by using either of the following methods:
   - Find the target cluster and click Add/Remove Node in the Actions column.

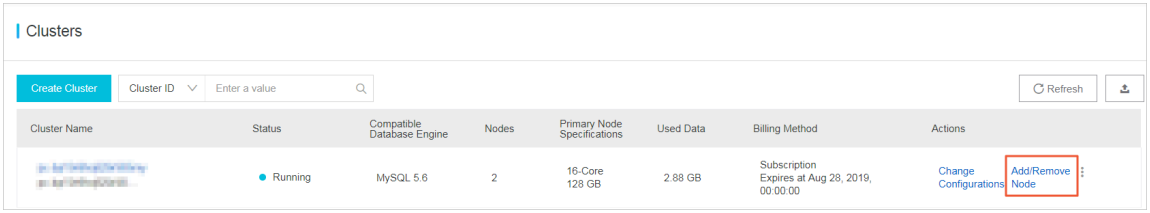

   - Find the target cluster, click the cluster ID, and then click Add/Remove Node in the Node Information section.

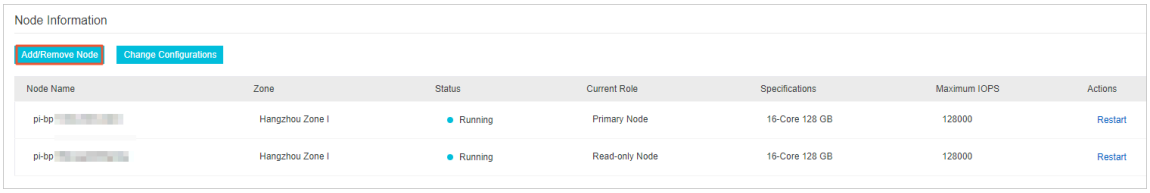

4. Select Remove Node and click OK.

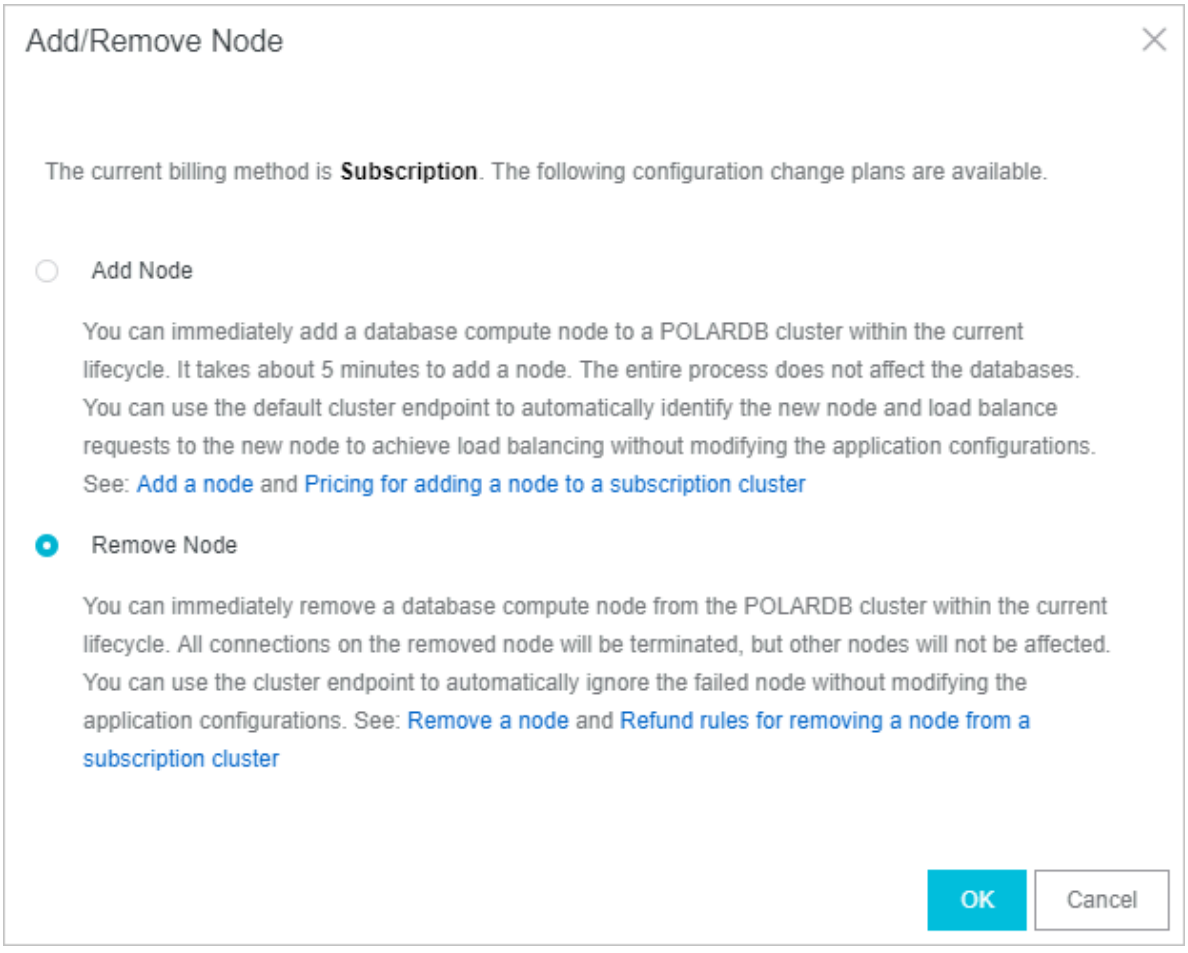

5. Click — next to the node that you want to remove. In the dialog box that appears, click OK.

> Note:
> You must keep at least one read-only node in the cluster to ensure high availability.

6. Read and agree to the service agreement by selecting the check box, and click OK.

Related API operations

API operation	Description
#unique_46	Adds a node to an ApsaraDB for POLARDB cluster.
#unique_47	Changes the specifications of a node in an ApsaraDB for POLARDB cluster.
#unique_48	Restarts a node in an ApsaraDB for POLARDB cluster.
#unique_49	Removes a node from an ApsaraDB for POLARDB cluster.

7.4 Change specifications

You can change the specifications of your cluster to meet business requirements. POLARDB supports capacity scaling in three dimensions:

- Scale up or down the computing capacity: Upgrade or downgrade the specifications of a cluster. [DO NOT TRANSLATE]
- Scale in or out the computing capacity: Add or delete read-only instances. For more information about the detailed procedures, see Add or delete read-only instances.
- Scale in or out the storage capacity: The storage capacity is provisioned in a serverless model. As your data increases in size, the storage is automatically expanded.

This topic describes how to upgrade or downgrade the specifications of a POLARDB cluster. It will take only 5-10 minutes for the new specification of each instance to take effect.

Note

- Specification upgrades or downgrades only apply to clusters. You cannot change the specifications of an instance.
- The specification upgrades or downgrades will not affect the existing data in the cluster.
- We recommend that you modify cluster specifications during your service off-peak periods. During a specification upgrade or downgrade, the POLARDB service will be disconnected for a few seconds and some of the functions will be disabled. You will need to reconnect from your applications once POLARDB is disconnected.
- You can only change cluster specifications when the cluster does not have pending specification changes..

Procedure

1. Log on to the POLARDB console.
2. Select a region.
3. Select Clusters, and find your targeted cluster. Click Morein the Actions column of the specific cluster, and then select Upgrade Cluster or Downgrade Cluster.
4. Select a specification.

   Note:
   All instances in a cluster have the same specifications.
5. Read and accept the Terms of Service, and click Pay Now.

   Note:
   It will take only 5-10 minutes for the new specification of each instance to take effect.

7.5 View clusters

View a list of clusters and the details of a cluster

1. Log on to the ApsaraDB for POLARDB console.

2. Select a region in the upper-left corner to view all the clusters that you want to deploy in this region.

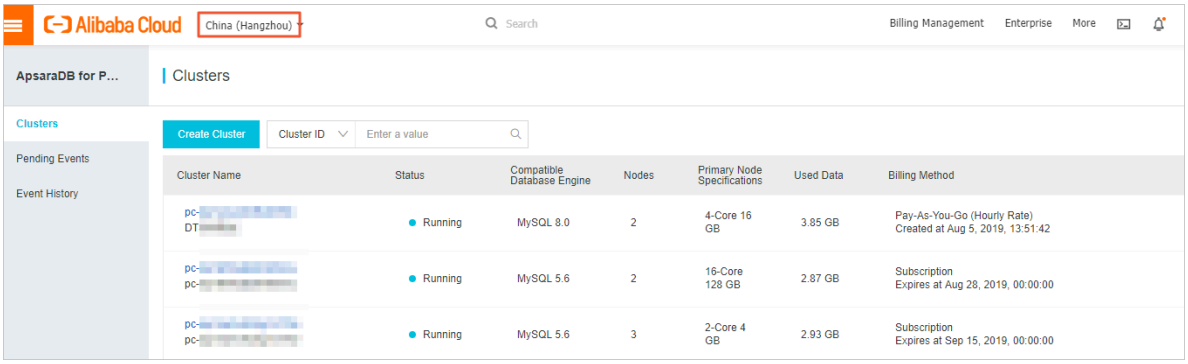

3. Click a cluster ID to go to the Basics page.

   You can view basic information, billing information, access information, and node information on the page.

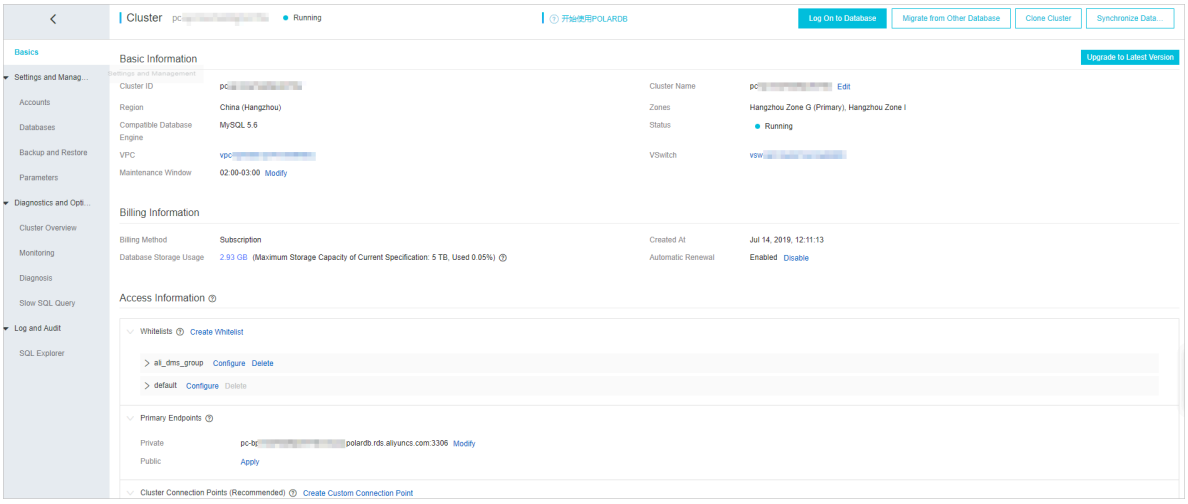

Related API operations

API operation	Description
#unique_23	Creates an ApsaraDB for POLARDB cluster.
#unique_38	Queries a list of ApsaraDB for POLARDB clusters.
#unique_39	Queries the detailed information of a specified ApsaraDB for POLARDB cluster.

7.6 Set cluster parameters

This topic describes how to modify parameter values of a cluster in the ApsaraDB for POLARDB console. For more information about the parameters, see Server System Variables.

Important notes

- You must modify parameter values according to the Value Range column on the Parameters page.

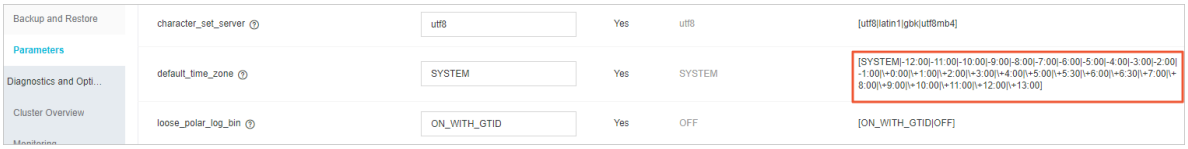

- For some parameters, you need to restart all nodes after the parameter values are modified. We recommend that you make appropriate service arrangements before you restart the nodes. Proceed with caution. You can determine whether the modification of a parameter value requires a node restart according to the value in the Force Restart column on the Parameters page.

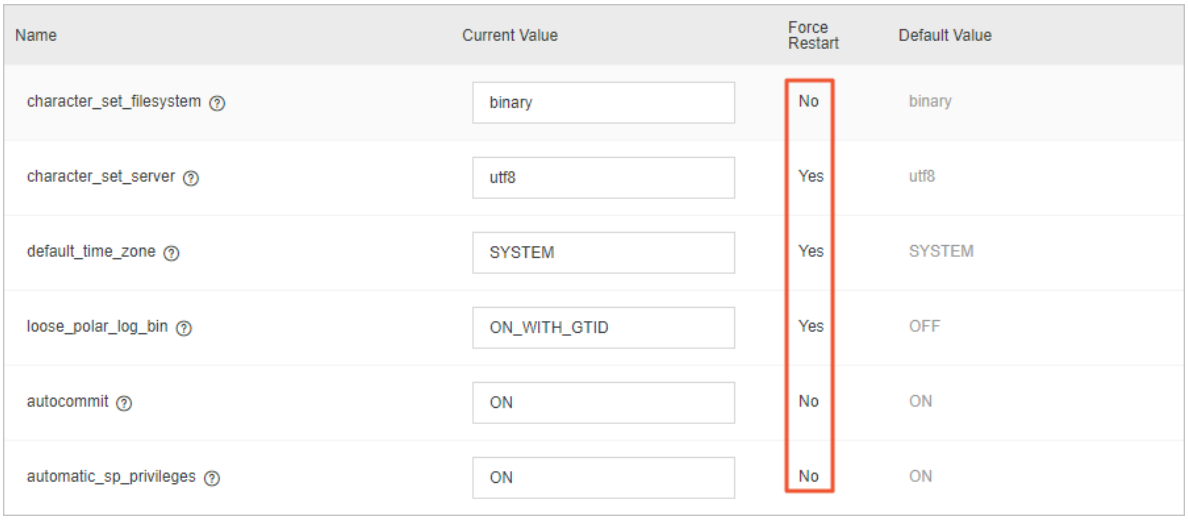

Procedure

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the left-side navigation pane, choose Settings and Management > Parameters.

5. Modify the values of one or more parameters in the Current Value column, and click Apply Changes.

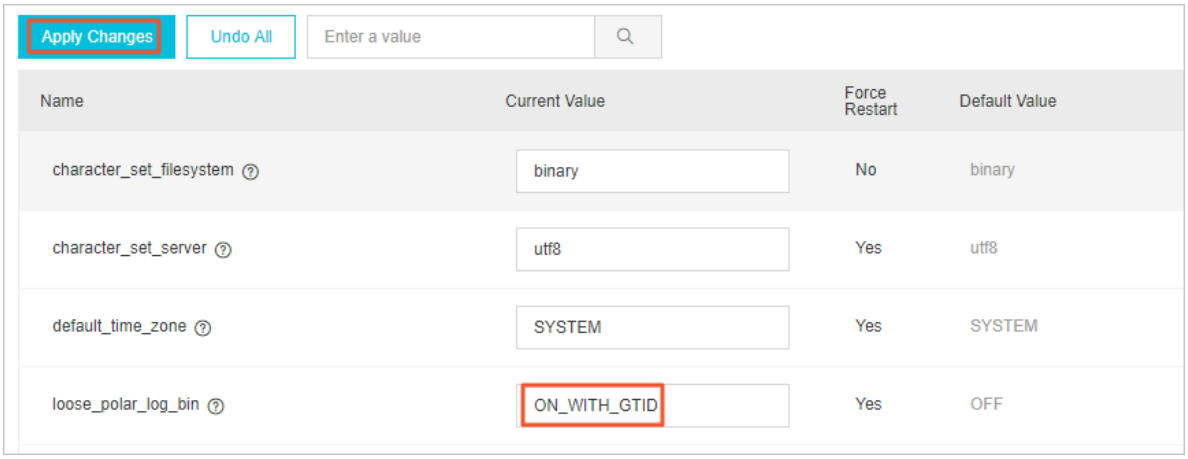

6. In the Save Changes dialog box that appears, click OK.

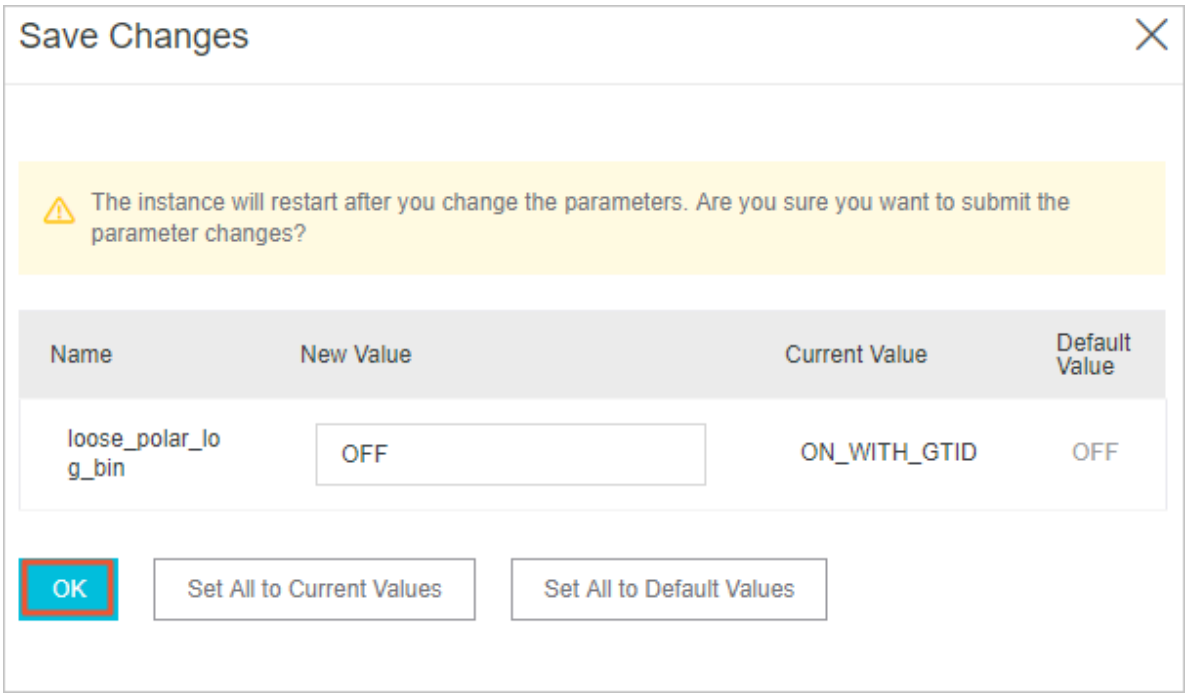

Related API operations

API operation	Description
#unique_54	Views cluster parameters.
#unique_55	Modifies the values of cluster parameters .

7.7 Set a maintenance window

To ensure the stability of ApsaraDB for POLARDB, the backend system performs maintenance operations on the clusters from time to time. We recommend that you set the maintenance window within the off-peak hours of your business to minimize the impact on the business during the maintenance process.

Important notes

- Before the maintenance is performed, ApsaraDB for POLARDB sends SMS messages and emails to contacts listed in your Alibaba Cloud account.
- To ensure stability during the maintenance process, clusters first enter the Under Maintenance status before the preset maintenance window arrives on the day of maintenance. When a cluster is in this status, normal data access to the database is not affected. However, except for the account management, database management, and IP address whitelisting functions, other services concerning changes (such as common operations like upgrade, degrade, and restart) are unavailable in the console of this cluster. Query services such as performance monitoring are still available.
- Within the maintenance window of a cluster, the cluster may experience one or two disconnections. Make sure that your application has an automatic reconnecti on mechanism. The cluster restores to the normal status immediately after the disconnection occurs.

Procedure

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the Basic Information section on the Basics page, click Modify next to Maintenance Window.

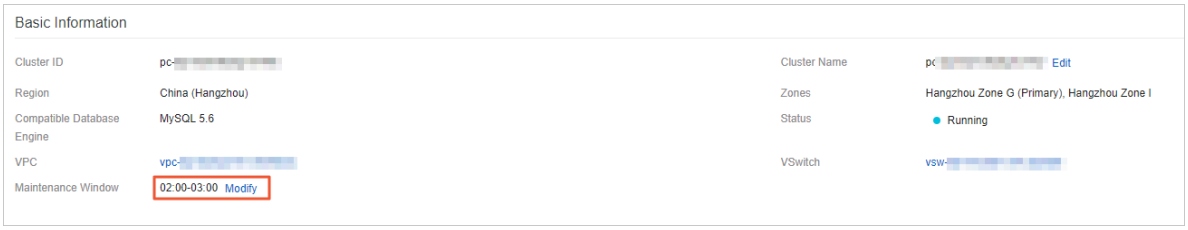

5. In the Modify Maintenance Window dialog box that appears, select a maintenance window for the cluster and click Submit.

Related API operations

API operation	Description
CreateDBCluster	Creates an ApsaraDB for POLARDB cluster.
ModifyDBClusterMaintainTime	Modifies the maintenance window for an ApsaraDB for POLARDB cluster.

7.8 Restart a node

You can manually restart a node when the number of connections exceeds the threshold or any performance issue occurs on the node. Restarting a node will cause service interruptions. We recommend that you make appropriate service arrangements before you restart the nodes. Proceed with caution.

Procedure

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the Node Information section on the Basics page, find the node to be restarted.
5. Click Restart in the Actions column of the node.

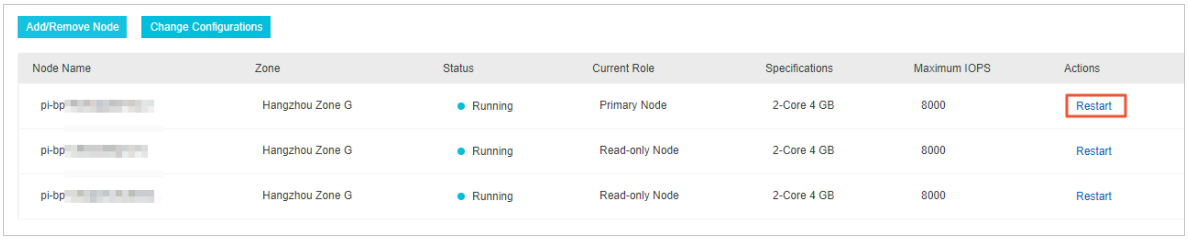

6. In the dialog box that appears, click OK.

Related API operations

API operation	Description
#unique_48	Restarts a database node.

7.9 Release a cluster

You can manually release a Pay-As-You-Go cluster according to your business requirements.

Note

- A Subscription cluster (billed annually or monthly) cannot be manually released, and will be automatically released once it expires.
- A Pay-As-You-Go cluster can only be manually released when it is in Running status .
- All data in your cluster will be deleted when it is released. Exercise caution.
- This function is used to release a cluster, including all instances in the specified cluster. To release a read-only instance, see Add or delete read-only instances.

Procedure

1. Log on to the POLARDB console.
2. Select the region where the cluster resides.
3. In the left-side navigation pane, select Clusters and find the target cluster. Click More in the Actions column, and then click Release.

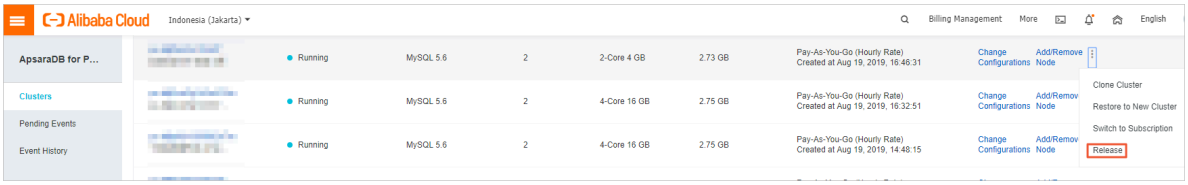

4. In the dialog box that appears, click OK.

8 Account Management

8.1 Overview

Console accounts

You can use the following accounts to log on to the console:

- Alibaba Cloud account: The account that allows flexible control of all your Alibaba Cloud resources and used for billing purposes. You must register an Alibaba Cloud account before purchasing any products.
- RAM user (optional): You can create and manage accounts in the Resource Access Management (RAM) console for resource sharing purposes. A RAM user does not own any resources, and is billed based on the corresponding Alibaba Cloud account.

Accounts for POLARDB cluster

You can use the following accounts to log on to your POLARDB cluster.

- Initial account: After purchasing a POLARDB cluster, you must create an initial account to access and manage the cluster. An initial account is an advanced user.
- Classic user: A classic user can be created and managed using SQL commands after you log on to the database with the initial account.

8.2 Register and log on to an Alibaba Cloud account

Register an Alibaba Cloud account

You can register an Alibaba Cloud account using the following two methods:

- On the Alibaba Cloud website (https://www.alibabacloud.com), click Free Account in the upper-right corner.

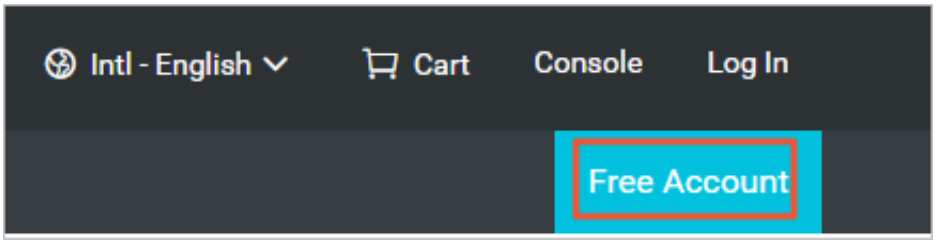

- Visit the Alibaba Cloud account registration page https://account.alibabacloud.com/register/intl_register.htm.

Log on to your Alibaba Cloud account.

Your Alibaba Cloud account and RAM user account have different logon pages.

- The logon page for Alibaba Cloud accounts is https://account.alibabacloud.com/login/login.htm.

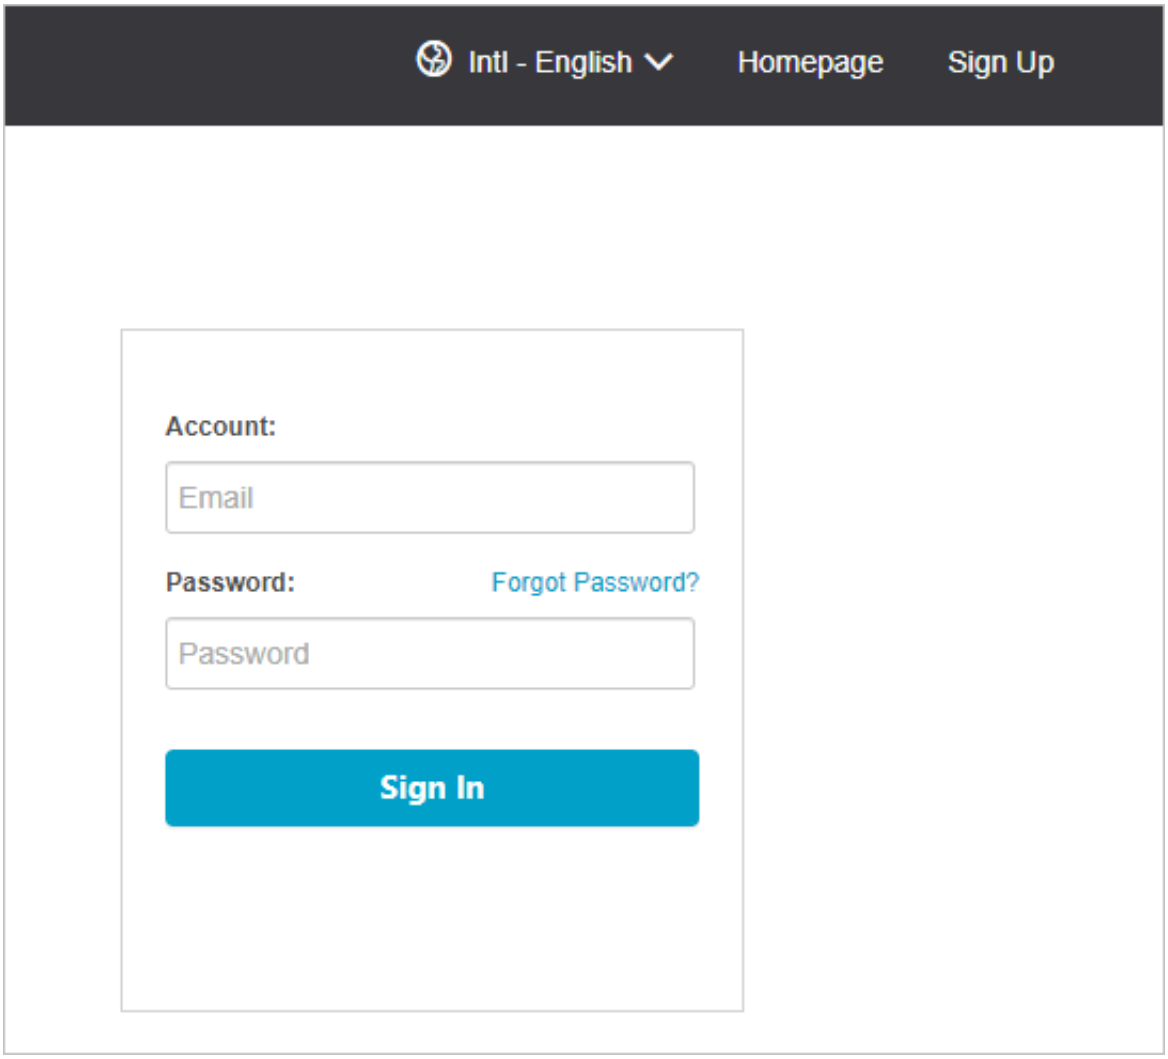

- The logon page for RAM users is https://signin-intl.aliyun.com/login.htm.

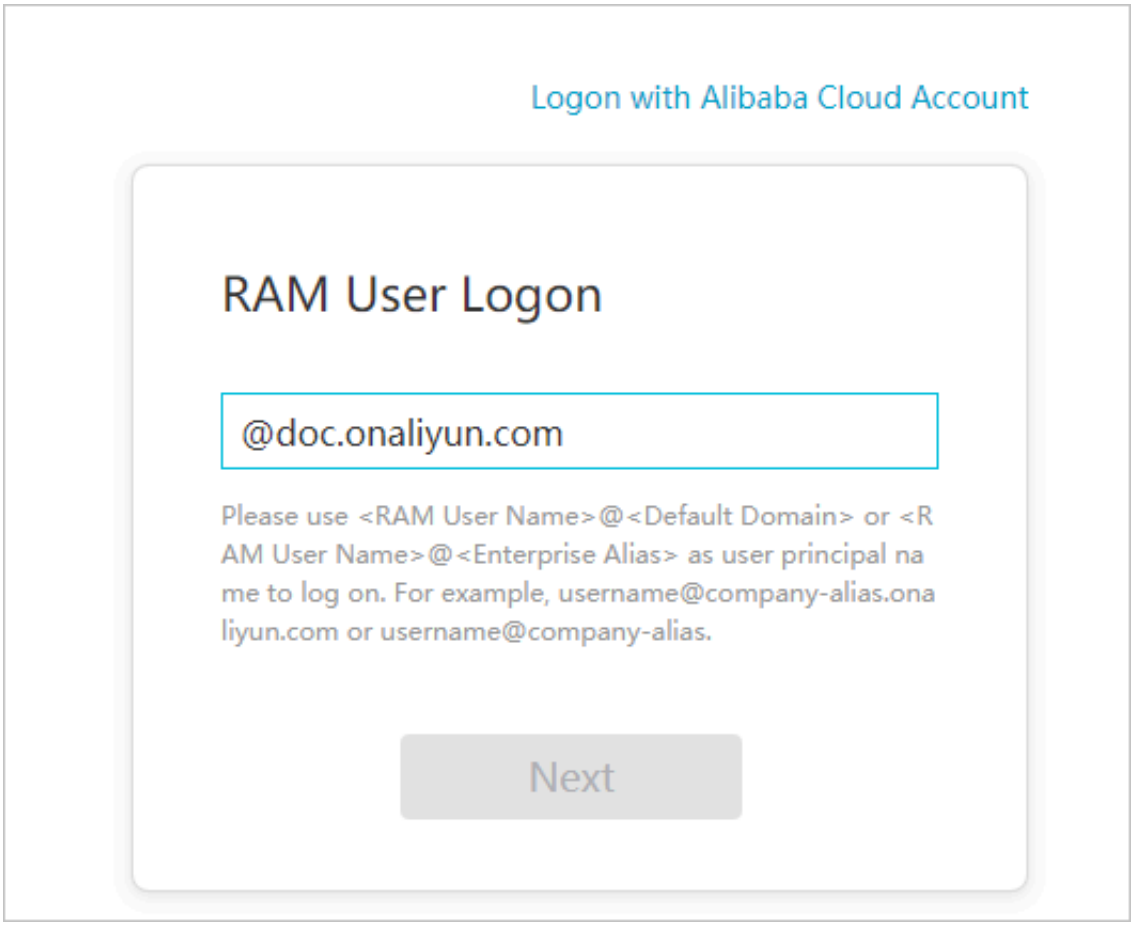

8.3 Create and authorize a RAM user

You can use your Alibaba Cloud account to access your ApsaraDB for POLARDB resources.

If you want to share the resources under your Alibaba Cloud account with other users, create and authorize a Resource Access Management (RAM) user. The RAM user can then be used to access specified resources.

Create a RAM user

1. You can use an Alibaba Cloud account or a RAM user to create one or more RAM users. First, log on to the RAM console.
   - Click here to log on with your Alibaba Cloud account.
   - Click here to log on with your RAM user.

   Note:

> Enter the RAM username in the format of `RAM username @ enterprise alias` on the logon page.

2. In the left-side navigation pane, click Identities, and click Users.
3. Click Create User.

> Note:
> To create multiple RAM users at a time, click Add User.

4. Specify the Logon Name and Display Name parameters.
5. In the Access Mode section, select Console Password Logon.
6. Under Console Password Logon, select Automatically Generate Default Password or Custom Logon Password.
7. Under Password Reset, select Required at Next Logon or Not Required.
8. Under Multi-factor Authentication, select Not Required.
9. Click OK.

Grant permission to a RAM user on the Grants page

1. In the left-side navigation pane, click Permissions, and click Grants.
2. Click Grant Permission.
3. Under Principle, enter the username, and click the target RAM user.
4. In the Policy Name column, select the target policies by clicking the corresponding rows.

> Note:
> You can click X in the section on the right side of the page to delete the selected policy.

5. Click OK.
6. Click Finished.

Grant permission to a RAM user on the Users page

1. In the left-side navigation pane, click Identities, and click Users.
2. In the User Logon Name/Display Name column, find the target RAM user.
3. Click Add Permissions. On the page that appears, the principle is automatically filled in.

4. In the Policy Name column, select the target policies by clicking the corresponding rows.

    Note:
   You can click X in the section on the right side of the page to delete the selected policy.

5. Click OK.
6. Click Finished.

Log on as a RAM user

Prerequisites: You must complete the preceding authorization procedures.

You can log on as a RAM user at the following addresses:

- Universal logon address: https://signin-intl.aliyun.com/login.htm.

  If you log on at the universal logon address, you must enter the RAM username and company alias manually. The address format is `RAM username @ company alias`.

- Dedicated logon address: You can view the logon address dedicated to your RAM users in the RAM console.

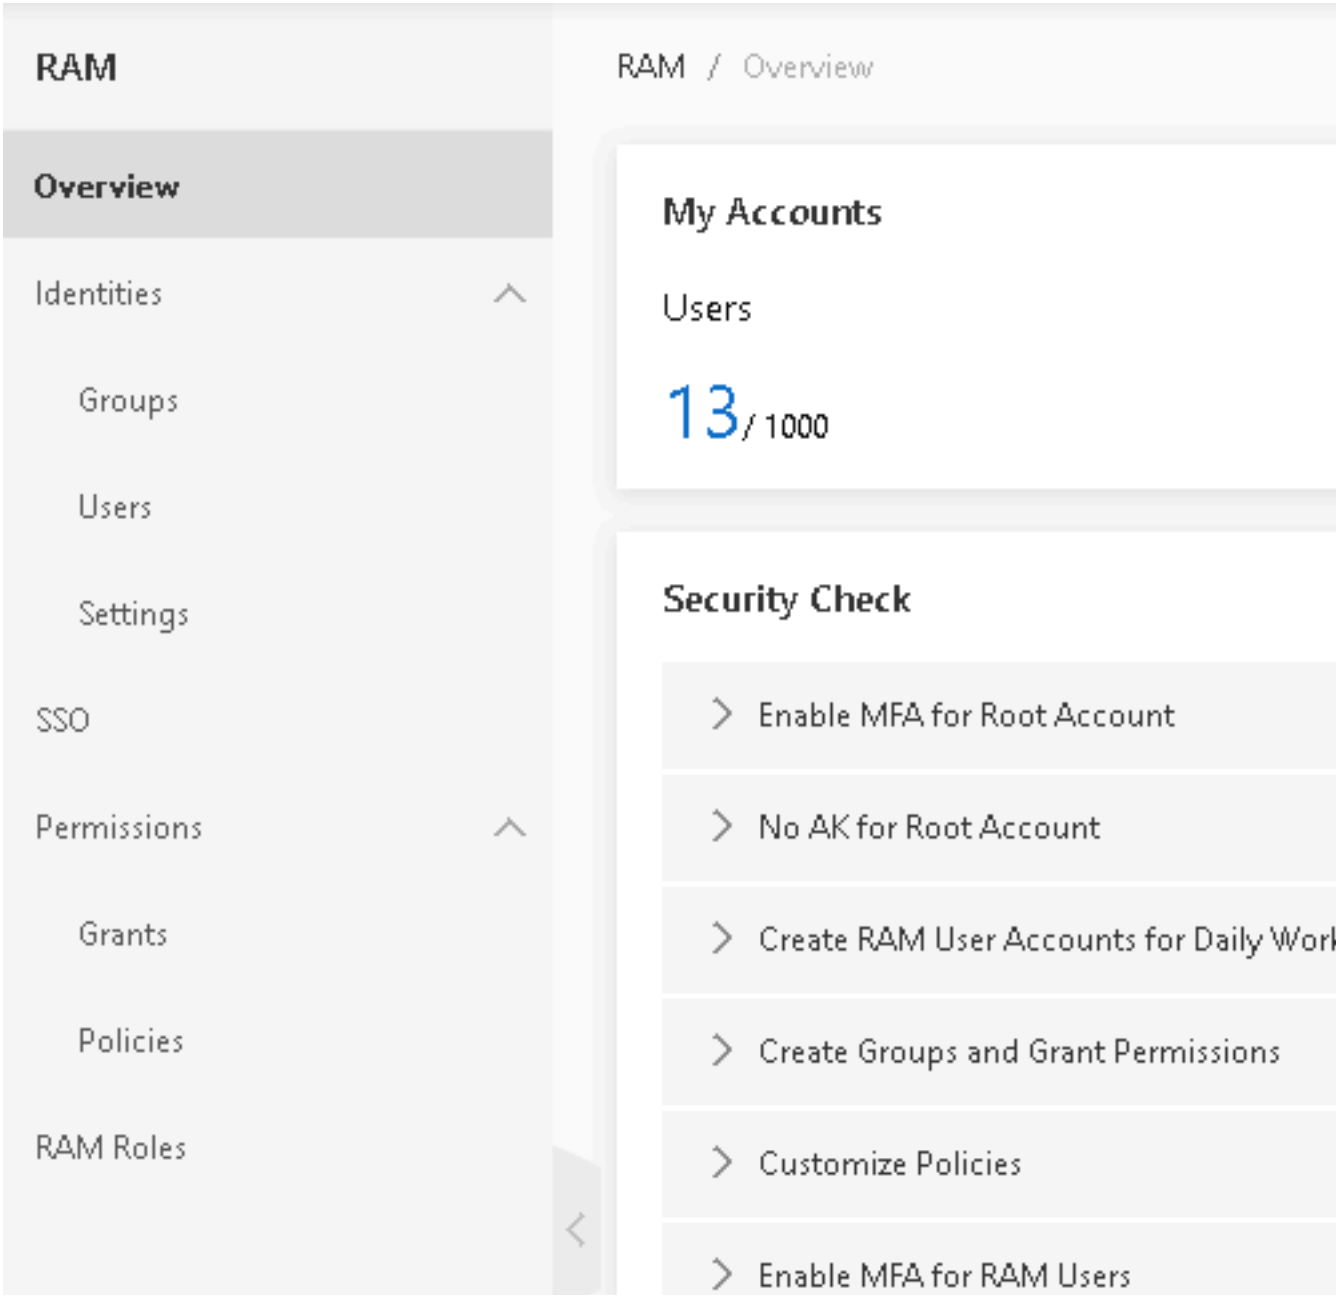

The system will enter your company alias automatically if you log on using this dedicated address. You only need to enter the RAM username.

More actions

You can also add a RAM user to a group, assign roles to a RAM user, and authorize a user group or roles. For more information, see RAM User Guide.

8.4 Create accounts for a POLARDB for MySQL cluster

POLARDB for MySQL supports two types of accounts: privileged accounts and standard accounts. You can manage all accounts in the POLARDB console.

Note:

For security purposes, POLARDB does not support the root accounts.

Account type	Description
Privileged accounts	· You can create and manage privileged accounts only in the POLARDB console. · Each POLARDB cluster can have only one privileged account. This privileged account has the permissions to manage all standard accounts and databases in the cluster. · More permissions are granted to a privileged account to achieve finer permission management such as granting the query permissions for tables by user. · The privileged account of a POLARDB cluster has the permissions for all databases in the cluster. · The privileged account of a POLARDB cluster can disconnect any account from the cluster.
Standard accounts	· You can create and manage standard accounts in the POLARDB console or by using SQL statements. · Each POLARDB cluster can have one or more standard accounts depending on the number of database cores. · You must manually grant the permissions for databases to standard accounts. · A standard account of a POLARDB cluster does not have the permissions to create or manage other accounts or disconnect other accounts from the cluster.

Create a privileged account

1. Log on to POLARDB console.
2. Find the target POLARDB cluster and click its ID.
3. In the left-side navigation pane, click Accounts.
4. Click Create Account.

5. Set the following parameters.

Parameter	Description
Account Name	Enter the name of the privileged account. The account name: · Must begin with a letter and end with a letter or digit. · Can contain lowercase letters, digits, and underscores (_). · Must be 2 to 16 characters in length. · Cannot be any reserved username, such as root and admin.
Account Type	Select Privileged Account. Note: If a privileged account already exists, you cannot select Privileged Account because each POLARDB cluster can have only one privileged account.
Password	Enter the password of the privileged account. The password: · Must contain at three of the following types of characters : uppercase letters, lowercase letters, digits, and special characters. · Must be 8 to 32 characters in length. · Must support only the following special characters: ! @ # $ % ^ & * ( ) _ + - =
Confirm Password	Re-enter the password.
Description	Enter the relevant account information for easy account management. The description: · Cannot start with http:// or https://. · Must start with a letter. · Can contain uppercase letters, lowercase letters, digits, underscores (_), and hyphens (-). · Must be 2 to 256 characters in length.

Create a standard account

1. Log on to the POLARDB console.
2. Find the target POLARDB cluster and click its ID.
3. In the left-side navigation pane, click Accounts.
4. Click Create Account.

5. Set the following parameters.

Parameter	Description
Account Name	Enter the name of the standard account. The account name: · Must begin with a letter and end with a letter or digit. · Can contain lowercase letters, digits, and underscores (_). · Must be 2 to 16 characters in length. · Cannot be any reserved username, such as root and admin.
Account Type	Select Standard Account.
Databases	Grant the permissions for one or more databases to the standard account. This parameter is optional. You can grant permissions to the standard account after you create it. a. Select one or more databases from the left section and click the right arrow to add them to the right section. b. In the right section, select Read&Write,ReadOnly, orDMLOnly permissions for each database.
Password	Enter the password of the standard account. The password: · Must contain at three of the following types of characters : uppercase letters, lowercase letters, digits, and special characters. · Must be 8 to 32 characters in length. · Must support only the following special characters: ! @ # $ % ^ & * ( ) _ + - =
Confirm Password	Re-enter the password.
Description	Enter the relevant account information for easy account management. The description: · Cannot start with http:// or https://. · Must start with a letter. · Can contain uppercase letters, lowercase letters, digits, underscores (_), and hyphens (-). · Must be 2 to 256 characters in length.

6. Click OK.

Reset the permissions of a privileged account

If the privileged account of a POLARDB cluster encounters a problem (for example, the permissions are revoked by mistake), you can enter the password of the privileged account to reset permissions.

1. Log on to the POLARDB console.
2. Find the target POLARDB cluster and click its ID.
3. In the left-side navigation pane, click Accounts.
4. Find the privileged account and in the Actions column click Reset Permissions.
5. In the displayed dialog box, enter the password of the privileged account and click OK.

APIs

API	Description
#unique_66	Used to create an account.
#unique_67	Used to list accounts.
#unique_68	Used to modify the description of an account.
#unique_69	Used to change the password of an account.
#unique_70	Used to grant permissions to an account.
#unique_71	Used to revoke the permissions of an account.
#unique_72	Used to reset the permissions of an account.

8.5 Manage the initial account of a POLARDB cluster

Note

An initial account cannot be deleted once created. You cannot change the account username, but you can change the password.

Create an initial account

See ../DNPOLA1840729/EN-US_TP_3016.dita#concept_ew4_wmq_tdb.

Reset password

1. Find the specified cluster or instance from the lists.
2. Click the cluster or instance ID, or click Manage in the Actions column of the target cluster or instance.
3. Click Change Password in the Access Information section.
4. In the dialog box that appears, enter a new password, and click OK.

9 Database management

You can create and manage all databases in the ApsaraDB for POLARDB console.

Create a database

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the `Cluster Name` column.
4. In the left-side navigation pane, choose Settings and Management > Databases.
5. Click Create Database.
6. In the dialog box that appears, set parameters for creating a database. The following table describes the parameters.

Parameter	Description
Database Name	· It must start with a letter and end with a letter or digit. · It can contain lowercase letters, digits, underscores (_), and hyphens (-). · It must be 2 to 64 characters in length. · Each database name in an instance must be unique.
Supported Character Set	Select utf8mb4, utf8, gbk, or latin1. You can also select other required character sets from the drop-down list on the right.
Authorized Account	Select the account that you want to authorize for accessing this database. You can leave this parameter blank, and bind an account after the database is created. Note: Only standard accounts are available in the drop-down list. The privileged account has all the permissions on all databases. You do not need to authorize the privileged account to access the database that you create.
Account Permission	Select the permission that you want to grant to your account. Valid values: Read and Write, Read Only, and DML Only.

Parameter	Description
Description	Enter the remarks of the database to facilitate subsequent database management. The requirements are as follows: · The description cannot start with http:// or https://. · The description must start with an uppercase or lowercase letter or a Chinese character. · The description can contain uppercase or lowercase letters, Chinese characters, digits, underscores (_), and hyphens (-). · The description must be 2 to 256 characters in length.

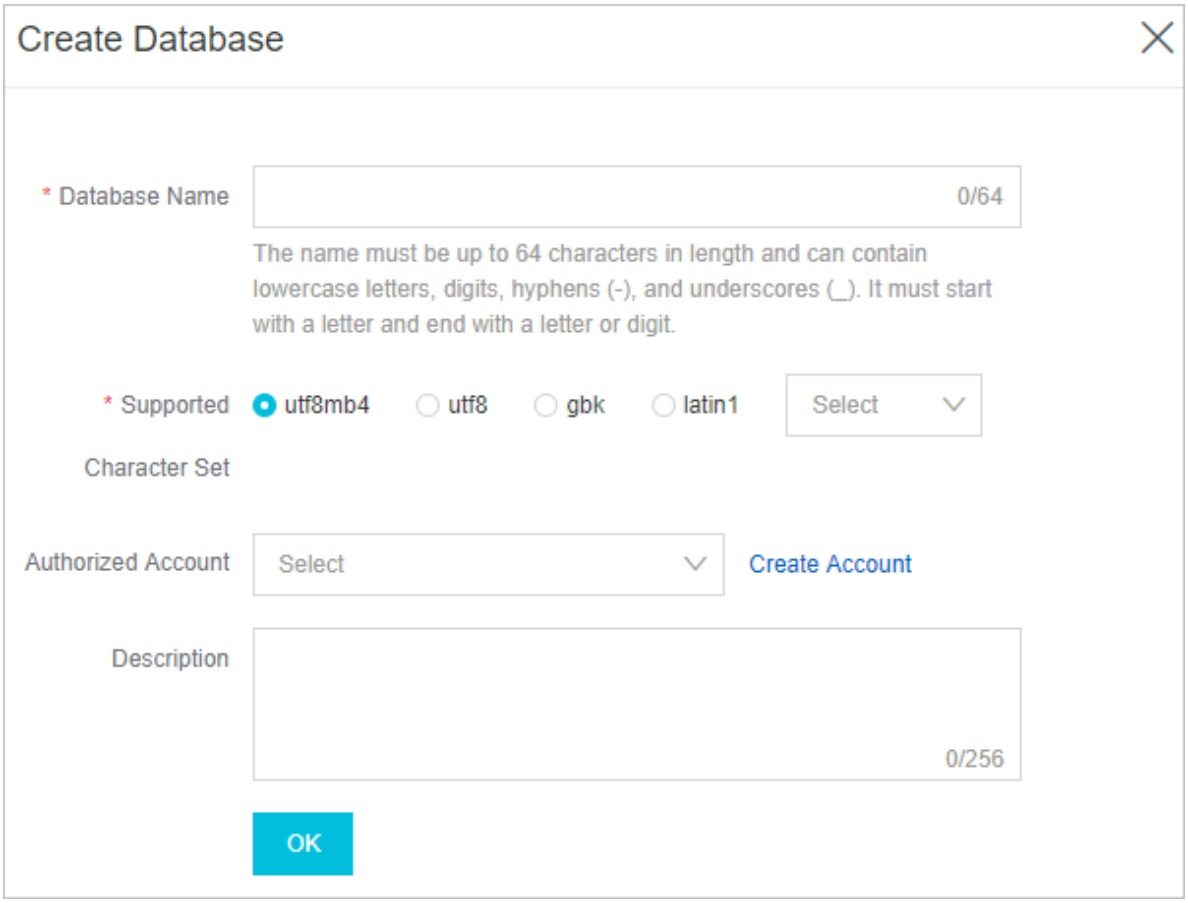

7. Click OK.

Delete a database

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the `Cluster Name` column.
4. In the left-side navigation pane, choose Settings and Management > Databases.
5. Find the target database and click Delete in the `Actions` column.
6. In the dialog box that appears, click OK.

Related API operations

API operation	Description
#unique_75	Creates a database.
#unique_76	Views the database list.
#unique_77	Modifies the description of a database.
#unique_78	Deletes a database.

10 Backup and restore

10.1 Back up data

ApsaraDB for POLARDB uses a physical backup (snapshot backup), which is automatically performed once a day. You can also manually start a backup. Both the automatic backup and manual backup do not affect the normal running of the cluster . Backup files are retained for seven days.

Backup types

Backup type	Description
Automatic backup	· It is performed once a day by default. You can configure the time period and cycle for an automatic backup. For more information, see Configure an automatic backup. · Backup files cannot be deleted.
Manual backup	· It can be started at any time. You can create a maximum of three manual backups for a cluster. For more information, see Create a manual backup. · Backup files can be deleted.

Pricing

Currently, the storage occupied by ApsaraDB for POLARDB backup files is free of charge.

Configure an automatic backup

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the left-side navigation pane, choose Settings and Management > Backup and Restore.

5. Click Backup Settings.

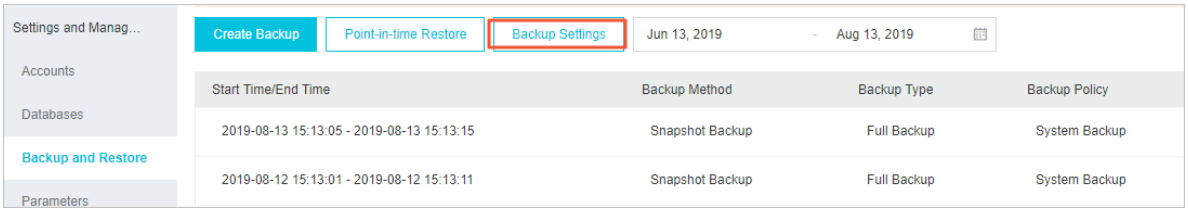

6. In the dialog box that appears, configure the time period and cycle for an automatic backup.

> Note:
> For security reasons, an automatic backup must be performed at least twice a week.

Create a manual backup

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the left-side navigation pane, choose Settings and Management > Backup and Restore.
5. Click Create Backup.

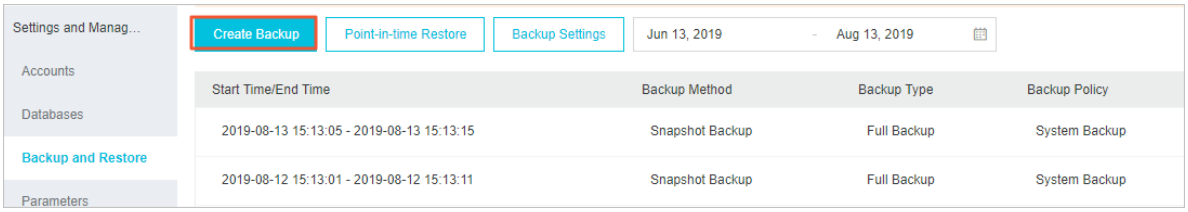

6. In the dialog box that appears, click OK.

> Note:
> You can create a maximum of three manual backups for a cluster.

Restore data

For more information, see Restore data.

10.2 Restore data

The process of restoring data of a POLARDB for MySQL cluster is as follows:

1. Restore historical data to a new cluster. You can choose either of the following methods to restore data:
   - Restore data to a specific point in time.
   - Restore data from a backup set (snapshot).
2. Log on to the new cluster and verify the data accuracy.

> Note:
> The restored cluster data contains the data and account information of the original cluster, excluding the parameter settings of the original cluster.

Restore data to a specific point in time

You can restore data to a specific point in time in the last seven days in a new cluster.

1. Log on to the ApsaraDB for POLARDB console.
2. Select the region where the original cluster resides.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the left-side navigation pane, choose Settings and Management > Backup and Restore.
5. Click Point-in-time Restore. In the dialog box that appears, click OK.

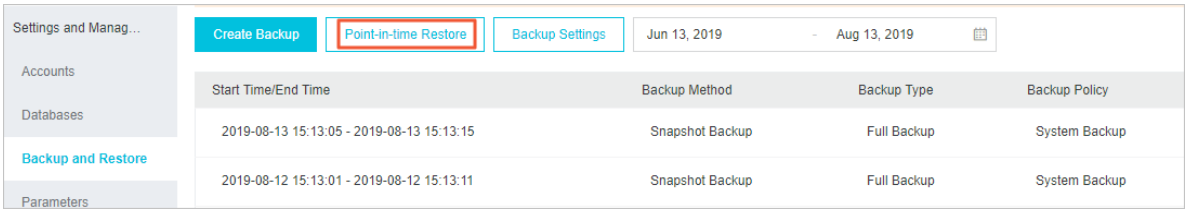

6. On the Clone Instance page, select a billing method for the new cluster:
   - Subscription: For the new cluster created, you need to pay the subscription fee for a compute cluster (with a primary node and a read-only node by default). The storage occupied by the new cluster is billed on an hourly basis based on the actual data volume. The payment will be deducted from your Alibaba Cloud account on an hourly basis. The subscription method is more cost-effective if you want to use the new cluster for a long term. You can save more with longer subscription periods.
   - Pay-As-You-Go (Hourly Rate): For the new cluster created, you do not need to pay any subscription fee for a compute cluster in advance. Use of the compute cluster is billed on an hourly basis. The storage occupied by the new cluster is billed on an hourly basis based on the actual data volume. The payment will be

deducted from your Alibaba Cloud account on an hourly basis. The pay-as-you-go method is suitable if you only want to use the new cluster for a short term. You can save the cost by releasing the cluster as soon as you complete the data restore.

7. Set the following parameters:
   - Clone Source Type: Select Backup Timepoint.
   - Backup Timepoint: Set it to a specific point in time in the last seven days.
   - Region: It is the same as the region of the original cluster. Use the default setting.
   - Primary Availability Zone: Use the default setting.
   - Network Type: Use the default setting.
   - VPC and Vswitch: We recommend that you use the default settings, namely, the VPC and VSwitch of the original cluster.
   - Database Engine: Use the default setting.
   - Node Specification: Clusters with different specifications have different storage capacity and performance. For more information, see #unique_37.
   - Number Nodes: Use the default setting. By default, the system will create a read-only node with the same specifications as the primary node.
   - Cluster Name: The system will automatically create a name for your POLARDB cluster if you leave it blank. You can rename the cluster after it is created.
   - Purchase Plan: Set this parameter if you create a cluster in subscription mode.
   - Number: The default value is 1, which cannot be modified.
8. Read and agree to the ApsaraDB for POLARDB service agreement by selecting the check box, and then complete the payment.

Restore data from a backup set (snapshot)

1. Log on to the ApsaraDB for POLARDB console.
2. Select the region where the original cluster resides.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the left-side navigation pane, choose Settings and Management > Backup and Restore.
5. Find the target backup set (snapshot) and click Restore in the Actions column. In the dialog box that appears, click OK.

6. On the page that appears, select a billing method for the new cluster:
   - Subscription: For the new cluster created, you need to pay the subscription fee for a compute cluster (with a primary node and a read-only node by default). The storage occupied by the new cluster is billed on an hourly basis based on the actual data volume. The payment will be deducted from your Alibaba Cloud account on an hourly basis. The subscription method is more cost-effective if you want to use the new cluster for a long term. You can save more with longer subscription periods.
   - Pay-As-You-Go: For the new cluster created, you do not need to pay any subscription fee for a compute cluster in advance. Use of the compute cluster is billed on an hourly basis. The storage occupied by the new cluster is billed on an hourly basis based on the actual data volume. The payment will be deducted from your Alibaba Cloud account on an hourly basis. The pay-as-you-go method is suitable if you only want to use the new cluster for a short term. You can save the cost by releasing the cluster as soon as you complete the data restore.
7. Set the following parameters:
   - Clone Source Type: Select Backup Set.
   - Clone Source Backup Set: Confirm that the backup set is the one that you want to restore from.
   - Region: It is the same as the region of the original cluster. Use the default setting.
   - Primary Availability Zone: Use the default setting.
   - Network Type: Use the default setting.
   - VPC and Vswitch: We recommend that you use the default settings, namely, the VPC and VSwitch of the original cluster.
   - Database Engine: Use the default setting.
   - Node Specification: Clusters with different specifications have different storage capacity and performance. For more information, see Node specifications.
   - Number Nodes: Use the default setting. By default, the system will create a read-only node with the same specifications as the primary node.
   - Cluster Name: The system will automatically create a name for your POLARDB cluster if you leave it blank. You can rename the cluster after it is created.
   - Purchase Plan: Set this parameter if you create a cluster in subscription mode.
   - Number: The default value is 1, which cannot be modified.

8. Read and agree to the ApsaraDB for POLARDB service agreement by selecting the check box, and then complete the payment.

FAQ

1. Q: Does the point-in-time restore method depend on binlogs? Is it possible to restore data to any point in time in the retention period of binlogs?

   A: The point-in-time restore method does not depend on binlogs. The cluster data can be restored to any point in time in the last seven days. The data restore is based on redo logs rather than binlogs.
2. Q: Is the data restore based on a full backup plus binlogs?

   A: The data restore is based on a full snapshot backup plus redo logs.

   The size of redo logs depends on your database write load. If a database is frequently written or updated, a large number of redo logs are generated. The system regularly uploads redo logs, and then clears the local redo logs. The local redo logs temporarily occupy the storage of the cluster and cost you a certain amount of fees. You will not be charged for the local redo logs after they are uploaded.

Related topics

#unique_89

11 Diagnostics and optimization

11.1 Diagnosis

12 Performance monitoring and alert configuration

The ApsaraDB for POLARDB console provides a variety of performance metrics for you to monitor the status of your instances.

Monitor performance

1. Log on to the ApsaraDB for POLARDB console.
2. Select a region.
3. Find the target cluster and click the cluster ID in the Cluster Name column.
4. In the left-side navigation pane, choose Diagnostics and Optimization > Monitoring.
5. You can view the performance information of a cluster or node according to your needs. For more information, see Metric description.
   - To monitor cluster performance, click the Cluster tab and set the monitoring time period.

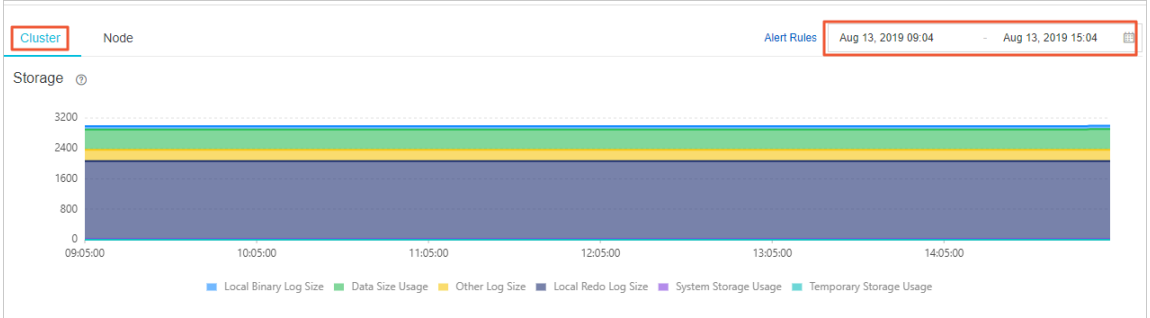

   - To monitor node performance, click the Node tab, select a node, and set the monitoring time period.

> Note:

You can click More at the bottom of the Node tab to view more metrics.

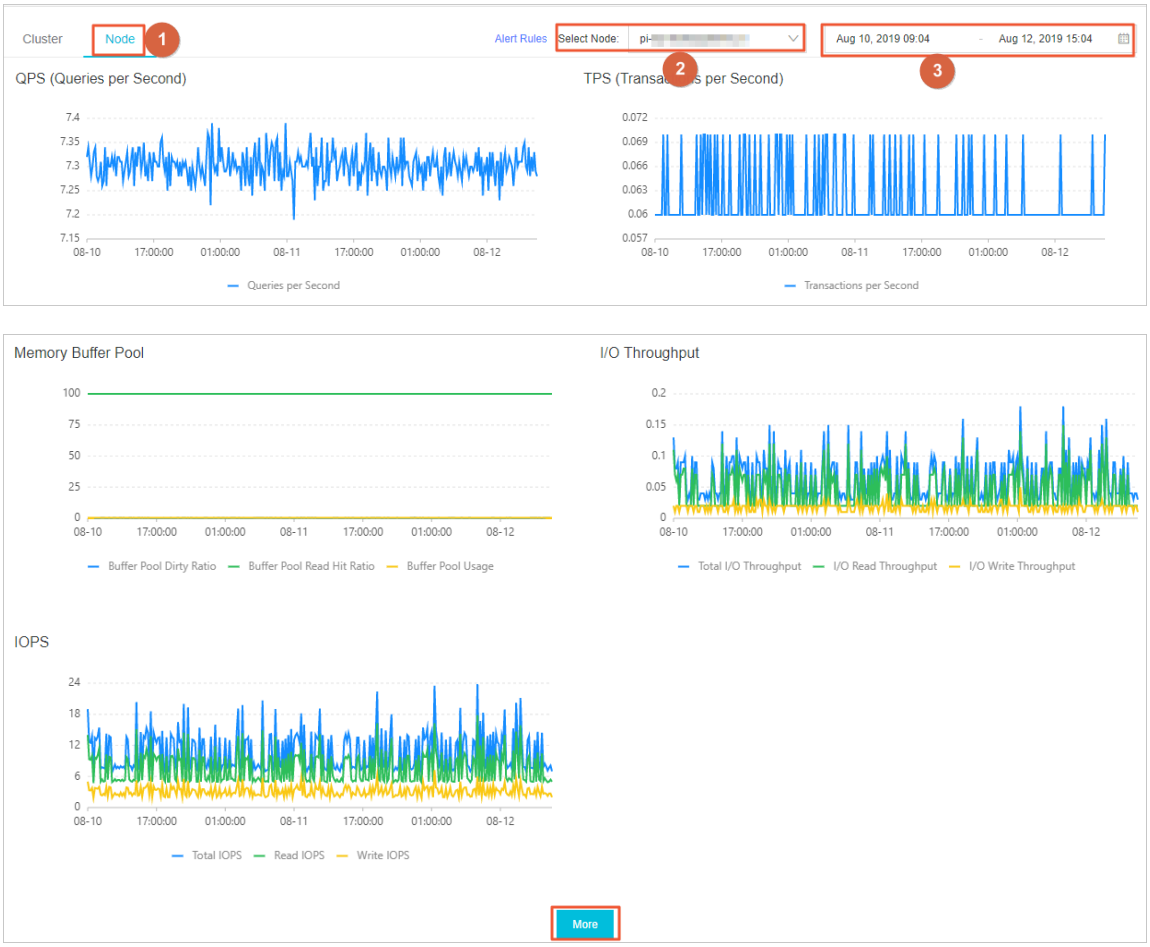

Metric description

Type	Metric	Description
Cluster	Storage	Displays the size of log files such as binlog and redolog, as well as the usage of data storage, system storage, and temporary storage.
	QPS	Displays the queries per second (QPS) of each node.
	TPS	Displays the transactions per second (TPS) of each node.
	CPU	Displays the central processing unit (CPU) usage of each node.
	Memory	Displays the memory usage of each node.
Node	QPS	Displays the QPS of the selected node.
	TPS	Displays the TPS of the selected node.
	CPU	Displays the CPU usage of the selected node.
	Memory	Displays the memory usage of the selected node.

Type	Metric	Description
	Connections	Displays the total number of connections and the number of active connections on the selected node.
	Operations	Displays the number of operations performed on the selected node per second, including the DELETE, INSERT, UPDATE, and REPLACE operations.
	Memory Buffer Pool	Displays the dirty ratio, read hit ratio, and usage of the buffer pool on the selected node.
	I/O Throughput	Displays the total I/O throughput, I/O read throughput, and I/O write throughput of the selected node.
	IOPS	Displays the input/output operations per second (IOPS) of the selected node, including the total IOPS, read IOPS , and write IOPS.
	Network	Displays the input and output traffic per second of the selected node.
	Scanned Rows	Displays the numbers of rows inserted, read, updated, and deleted per second on the selected node.
	InnoDB Read and Written Data	Displays the amount of data read from or written into the storage engine per second on the selected node.
	InnoDB Buffer Pool Requests	Displays the numbers of read and write operations performed on the buffer pool of the selected node per second.
	InnoDB Log Writes	Displays the number of log write requests per second and the number of times that data is synchronized to disks per second on the selected node.
	Temporary Table	Displays the number of temporary tables created per second on the selected node.

13 Clone a cluster

You can create a ApsaraDB for POLARDB cluster the same as an existing ApsaraDB for POLARDB cluster by cloning the data of the existing one. The data includes the account information, but excludes parameter settings of the cluster.

The data generated before the execution of the clone action is cloned. When cloning starts, the newly written data will not be cloned.

Procedure

1. Log on to the ApsaraDB for POLARDB console.
2. Select the region where the target cluster is located.
3. Find the cluster you want to clone. In the `Actions` column of the cluster, click the More icon, and then select Restore to New Cluster.
4. On the page that appears, set the parameters. The following table describes the parameters.

Parameter	Description
Clone Source Type	The type of the clone source. Select Current Cluster.
Region	The region where the cluster resides. The region of the new cluster is the same as that of the source cluster and cannot be modified.
Primary Availability Zone	· The zone of the new cluster. A zone is an independent physical area located within a region. There are no substantive differences between the zones. · You can deploy the ApsaraDB for POLARDB cluster and ECS instance in the same zone or in different zones.
Network Type	· The type of the network. Use the default setting. · ApsaraDB for POLARDB supports Virtual Private Cloud (VPC) networks only. A VPC is an isolated virtual network with higher security and performance than a classic network.

Parameter	Description
VPC Vswitch	The VPC and VSwitch of the new cluster. Select a VPC and a VSwitch from the corresponding drop-down lists, or create a VPC and a VSwitch. Note: Make sure that you place your ApsaraDB for POLARDB cluster and the ECS instance to be connected in the same VPC. Otherwise, they cannot intercommunicate through the internal network and achieve optimal performance.
Database Engine	The database engine of the new cluster. Use the default setting.
Node Specification	The node specification of the new cluster. Select a specification according to your needs. Clusters with different specifications have different storage capacity and performance. For more information, see #unique_37.
Number Nodes	The number of nodes in the new cluster. Use the default setting. By default, the system creates a read-only node with the same specification as the primary node.
Cluster Name	· Optional. The name of the new cluster. · The system will automatically create a name for your ApsaraDB for POLARDB cluster if you leave it blank. You can rename the cluster after it is created.
Purchase Plan	The subscription duration of the new cluster. This parameter is valid only for subscription clusters.
Number	The number of clusters. The default value 1 is used and cannot be modified.

5. Read the ApsaraDB for POLARDB Agreement of Service, select the check box to agree to it, and then complete the payment.

14 Enable binlogging

ApsaraDB for POLARDB is a cloud native database fully compatible with MySQL. By default, it uses more advanced physical logs instead of binlogs. However, to better integrate with the MySQL ecosystem, ApsaraDB for POLARDB allows you to enable binlogging. When binlogging is enabled, you can connect to the data products such as ElasticSearch and AnalyticDB. You can also synchronize data from POLARDB to RDS, from RDS to POLARDB, or between POLARDB clusters in a real time manner.

Prerequisites

The cluster was created after April 5, 2019. If the cluster was created on April 5, 2019 or earlier, you need to open a ticket to perform minor version upgrade. After that, you can enable binlogging in the console.

Pricing

The space used to store binlogs is a part of the cluster storage space. It is charged based on the pricing policies.

Precautions

- By default, the binlog files are saved for two weeks after binlogging is enabled. The binlog files that are generated more than two weeks ago are automatically deleted. You can modify the loose_expire_logs_hours parameter to set the duration for storing binlog files. The value ranges from 0 to 2376, in hours. The value 0 indicates that the binlog files are not automatically deleted.
- By default, the binlogging feature is disabled. To enable this feature, you need to restart the instance, which will cause service interruptions. We recommend that you arrange services appropriately before you restart an instance.
- After binlogging is enabled, the write performance is deteriorated, while the read performance is not affected.
- The primary endpoint directly points to the primary node that generates binlog files, ensuring higher compatibility and stability. We recommend that you use the primary endpoint of ApsaraDB for POLARDB when you pull, subscribe to, or

synchronize binlog files by using a tool such as DTS. You can view the primary endpoint on the Basic Information page, as shown in the following figure.

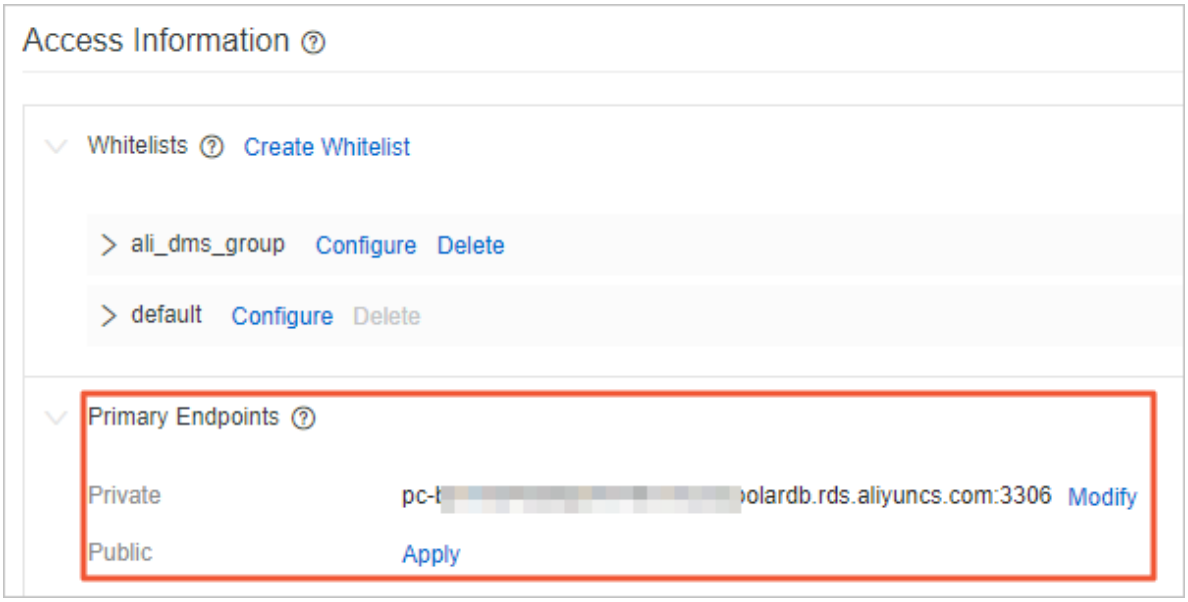

Enable binlogging

1. Log on to the ApsaraDB for POLARDB console.
2. Select the region where the target cluster is located.
3. Find the target cluster, and then click the cluster ID in the Cluster Name column.
4. In the left-side navigation pane, choose Settings and Management > Parameters.
5. Search for the loose_polar_log_bin parameter, change the value of the parameter to `ON_WITH_GTID`, and then click Apply Changes.

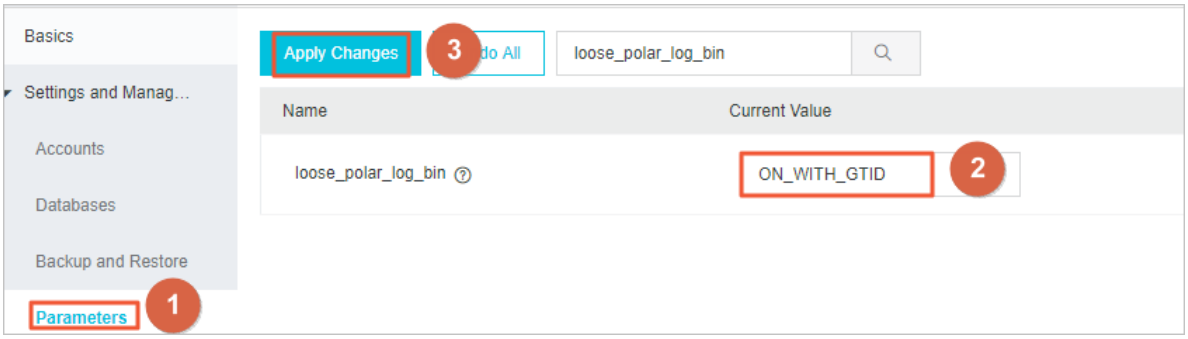

6. In the dialog box that appears, click OK.

> Note:
> If the error message Custins minor version does not support current action is displayed, open a ticket to enable binlogging.

FAQs

- **Q: How long can binlog files be stored?**

  A: By default, the binlog files are saved for two weeks after binlogging is enabled. The binlog files that are generated more than two weeks ago are automatically deleted. You can modify the loose_expire_logs_hours parameter to set the duration for storing binlog files. The value ranges from 0 to 2376, in hours. The value 0 indicates that the binlog files are not automatically deleted.

- **Q: How do I disable binlogging after it is enabled?**

  A: Set the loose_polar_log_bin parameter to `OFF`, and then submit the changes. The existing binlog files will not be deleted after binlogging is disabled.

- **Q: What is the impact of enabling binlogging on performance?**

  A: According to the test data, with 64 concurrent calls, there will be a write performance deterioration of 30% to 40% after binlogging is enabled. The write performance is optimized with the increase of concurrent calls, and will be continuously optimized in the future. The read performance will not be affected . For business scenarios with more read operations than write operations, the impact on the overall database performance is slight. For example, if the read-write ratio is 4:1, the overall performance is deteriorated by about 10%.

15 FAQs

Q: Why does the database still occupy a large amount of storage after it is deleted?

A: This is because the redo logs occupy storage. Usually, the logs occupy around 2 to 11 GB. They can occupy up to 11 GB: 8 GB (the 8 redo logs in the buffer pool) + 1 GB (the redo log being written) + 1 GB (the redo log created in advance) + 1 GB (the last redo log).

The number of redo log files in the buffer pool is determined by the loose_innodb_polar_log_file_max_reuse parameter. The default value is 8. You can modify this parameter to reduce the storage usage of redo logs. However, when a large amount of storage is occupied, performance may fluctuate slightly in a periodic manner.

Note:
After you modify the loose_innodb_polar_log_file_max_reuse parameter, the buffer pool will not be cleared immediately. It will be gradually released as Data Manipulation Language (DML) operations are performed. If you need to clear the buffer pool immediately, contact after-sales service representatives.

Q: What can I do if the disk space cannot be specified?

A: The disk space does not need to be manually specified. The system automatically scales the space according to the data volume.

ApsaraDB for POLARDB uses a storage cluster at the underlying layer to dynamicall y scale up the disk space without service interruption. When the disk space usage reaches 70%, the system automatically scales up the disk space without stopping the instance. Through this mechanism, the storage space of ApsaraDB for POLARDB can be billed based on usage.

Q: How does read-write splitting guarantee the read consistency?

A: The read-write splitting link records the log sequence number (LSN). Read requests are sent to read-only nodes that meet the requirements of the LSN. For more information, see Read-write splitting.

Q: How do I realize read-write splitting for ApsaraDB for POLARDB?

A: Use cluster connection points in the application to realize read-write splitting based on the configured reader nodes. You can also use custom cluster connection points.

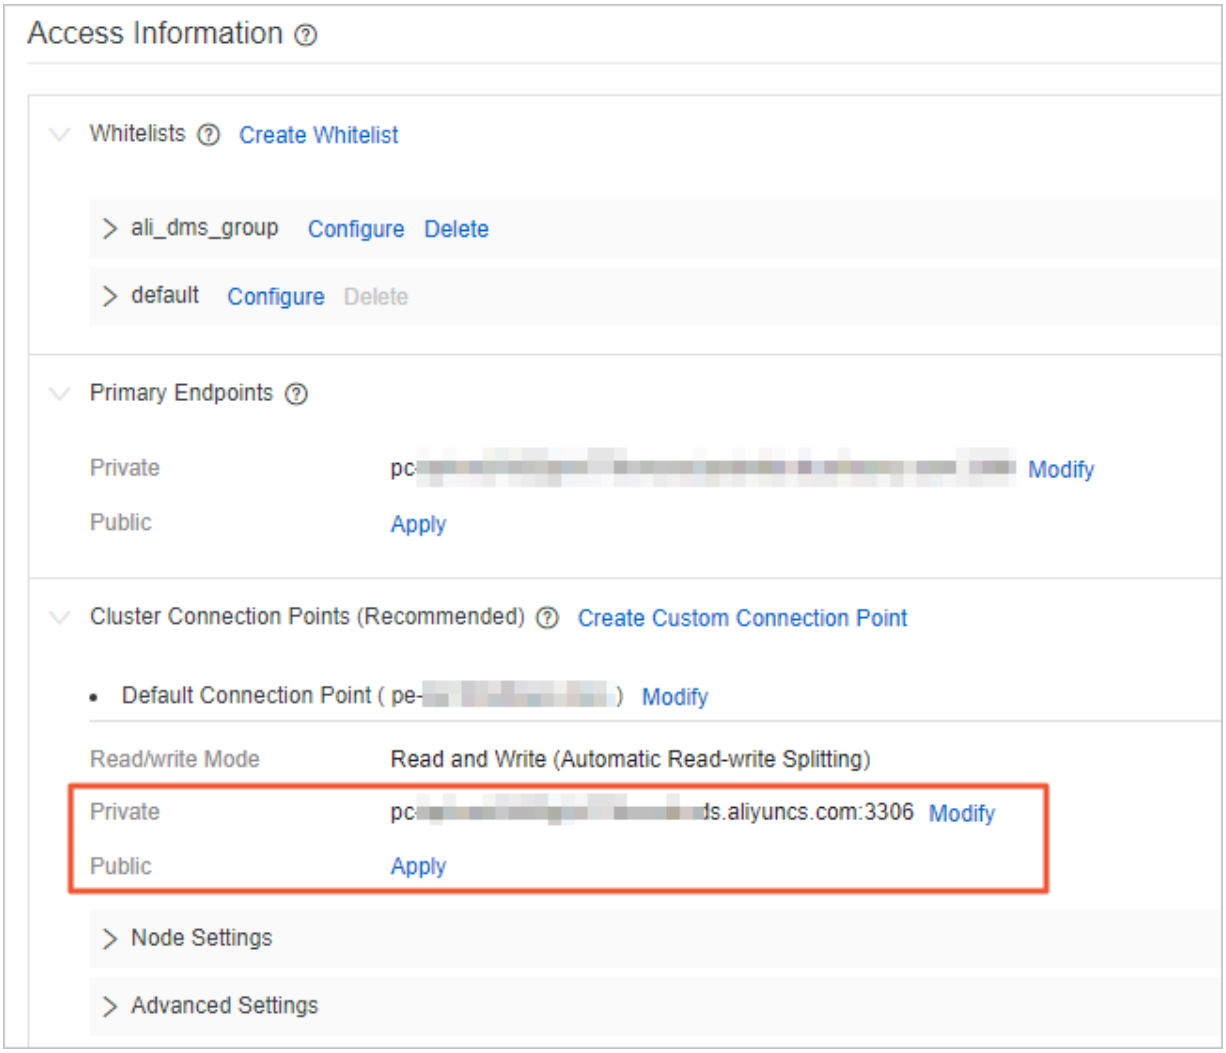

Q: If there are multiple read-only nodes, how do I set the specified Elastic Compute Service (ECS) instance to access the specified read-only nodes?

A: Set a custom cluster connection point, select the read-only nodes to be connected, and then use the custom cluster connection point on the specified ECS instance.

Q: Read-only nodes are loaded when I only use the primary endpoint. Does the primary endpoint support read-write splitting?

A: The primary endpoint does not support read-write splitting. It is always connected to the primary node. It is normal that the query per second (QPS) of read-only nodes is at a low level, which is not related to the primary endpoint.

Q: How do I find a slow SQL query?

A: After #unique_98, run the `show processlis t ;` command to find the SQL queries that have been running for a long period of time.

Q: How do I terminate a slow SQL query?

A: When a slow SQL query is found, view its ID and run the `kill < Id >` command to terminate the slow SQL query.

16 Appendix