ALIBABA CLOUD

Alibaba Cloud

ApsaraDB for RDS

Best Practices

Document Version: 20210111

Alibaba Cloud

Legal disclaimer

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

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

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

Table of Contents

1.Overview of best practices

This topic lists the best practices of ApsaraDB RDS in various business scenarios.

Database engine	Topic
All	• Build a high availability architecture
MySQL	• Best practices of X-Engine • Import data from Excel to ApsaraDB RDS for MySQL • Configure a cyclic event on an ApsaraDB RDS for MySQL instance • Select and create an optimal index for faster data access
SQL Server	• Connect Kingdee K/3 WISE to ApsaraDB RDS for SQL Server • Use SSRS for an ApsaraDB RDS SQL Server instance
PostgreSQL	• CREATE DATABASE usage instructions • Set the localized information and string sorting rules for a database • Functions and usage of PostgreSQL UPSERT • Update, delete, and insert data in batches • Locate the SQL statements with the greatest resource consumption • Real-time precision marketing (user selection) • Image recognition, face recognition, similarity-based retrieval, and similarity-based audience spotting

2.Build a high availability architecture

ApsaraDB for RDS provides a complete suite of high availability features such as the dedicated instance family, high availability-centered RDS editions, multi-zone deployment, and cross-region backup and restoration.

RDS editions and instance families

When you create an RDS instance, note the following high availability-related options:

- **Edition**: We recommend that you select the **High-availability**, **Enterprise**, or **Cluster (AlwaysOn)** Edition.
  - **High-availability**: The database system works in the classic high-availability architecture and consists of one primary instance and one secondary instance.
  - **Enterprise**: The database system consists of one primary instance and two secondary instances. The primary and secondary instances reside in three different zones within the same region to provide financial-level reliability.
  - **Cluster (AlwaysOn)**: This edition is only supported for SQL Server. The database system consists of one primary instance, one secondary instance, and up to seven read-only instances used to scale out the read capability.

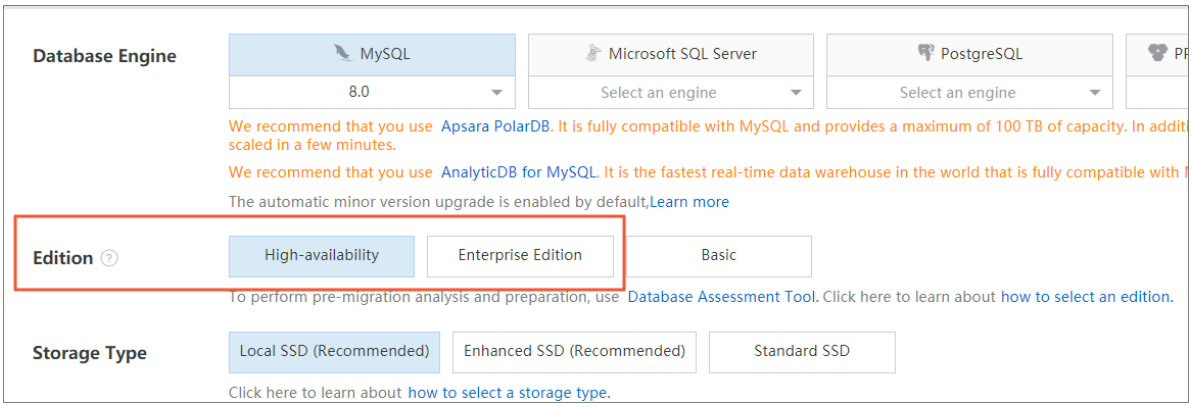

- **Zone**: ApsaraDB for RDS supports both single-zone deployment and multi-zone deployment. We recommend that you use multi-zone deployment. If your database system spans multiple zones, it can provide zone-level disaster recovery.
- **Instance Type**: We recommend that you select the **Dedicated Instance** or **Dedicated Host** family.
  - **Dedicated Instance**: A dedicated instance occupies the exclusive CPU and memory resources allocated to it. Its performance and stability are independent of the other instances deployed on the same physical host.
  - **Dedicated Host**: This is the top configuration of the **Dedicated Instance** family. A dedicated host instance occupies all resources on the physical host where it is housed.

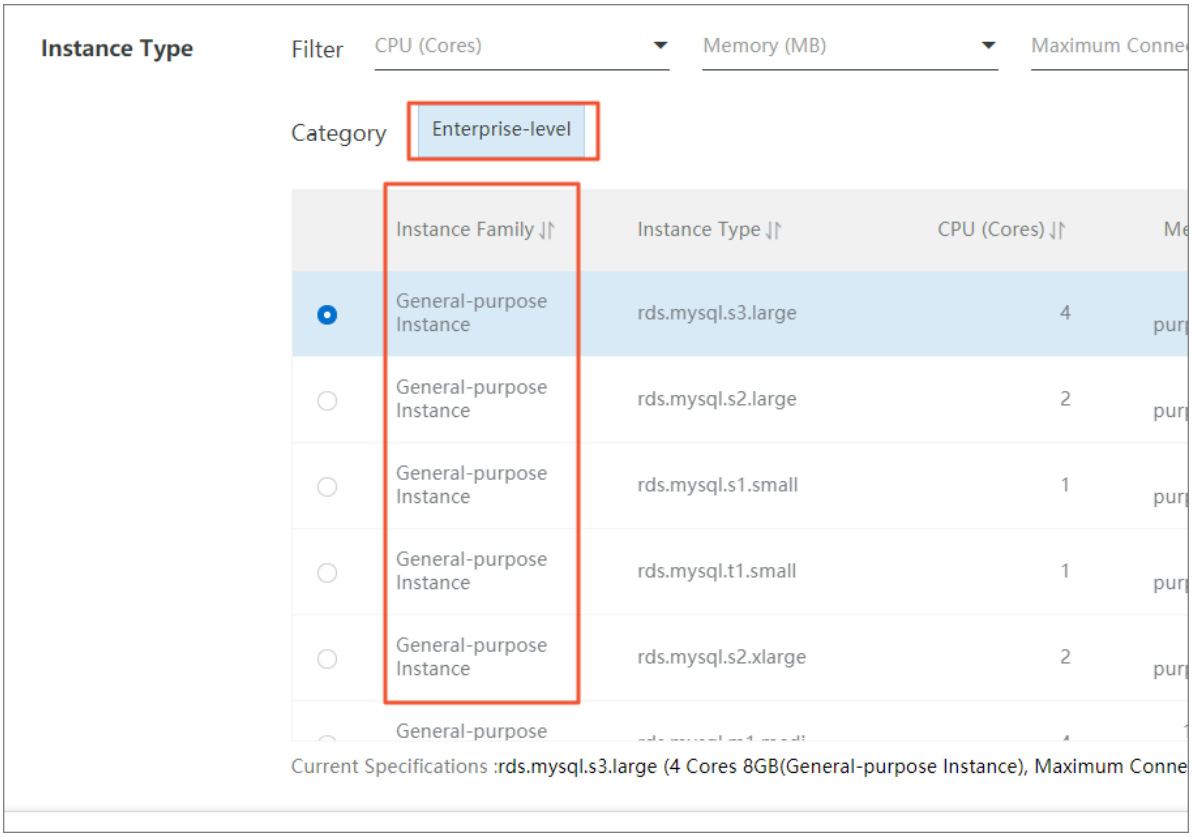

Automatic backup

We recommend that you configure an automatic backup policy for your ApsaraDB for RDS instance. If your instance becomes unavailable due to misoperations or other exceptions, you can use the backups to restore the instance to its latest state.

Cross-region disaster recovery

The cross-region disaster recovery feature helps secure your data and increases the availability of your RDS instances.

- Create a disaster recovery ApsaraDB RDS for MySQL instance: The primary instance and its remote disaster recovery instance synchronize data with each other in real time by using Data Transmission Service (DTS). Both the primary and disaster recovery instances are deployed based on the primary/secondary high availability architecture. If your application cannot connect to either the primary or secondary instance due to a natural disaster, you can switch services over to the remote disaster recovery instance and then update the endpoints on your application. This minimizes the downtime of your database system.
- Back up an ApsaraDB RDS for MySQL instance across regions: Your database system automatically replicates its backup files to OSS buckets in a different region.

Monitoring and alerting

To prevent unavailability caused by CPU, disk, memory, connection, or other exceptions, we recommend that you monitor and configure thresholds for the performance metrics of your RDS instances. If the value of a metric reaches the preset threshold, the system reports alerts. For more information, see Configure alert rules for an ApsaraDB RDS MySQL instance.

Data restoration

If you have built a high availability architecture for your database system as instructed above, your business can run without downtime under normal scenarios and can even be restored within a short time in the event of exceptions.

- If a single RDS instance of your database system is faulty, you can switch services over to another RDS instance. This operation is unavailable in the Basic Edition.
- In the multi-zone deployment solution, the primary instance can be switched to another zone if the current zone is faulty. In the single-zone deployment solution, you must wait until the fault is rectified or switch services over to the disaster recovery instance.
- If the region of your database system is faulty, you can switch services over to the disaster recovery instance. You also have the option to restore the data to a new RDS instance by using cross-region backup.

For more information about how to restore data, see the following topics:

- Restore the data of an RDS instance
- Restore individual databases and tables for an ApsaraDB RDS for MySQL instance
- Restore the data of an ApsaraDB RDS for MySQL instance across regions

3.Use DMS to import a logical backup file into an ApsaraDB for RDS instance

This topic describes how to import a logical backup file into an ApsaraDB for RDS instance by using Alibaba Cloud Data Management (DMS).
For more information, see Import data.

4.Authorize a RAM user to manage ApsaraDB for RDS instances

This topic describes how to authorize a RAM user to manage ApsaraDB for RDS instances by using Resource Access Management (RAM).

Prerequisites

A RAM user is created. For more information, see Create a RAM user.

Context

The essence of authorizing a RAM is to grant the RAM user the permissions to call API operations. For example, if you grant the RAM user the permission to call the CreateDBInstance API operation, the RAM user can create an RDS instance in the ApsaraDB for RDS console.

The following procedure shows how to grant a RAM user the permission to view RDS instances. The procedures to grant other permissions to a RAM user are similar.

Procedure

1. Log on to the RAM console.
2. In the left-side navigation pane, choose **Permissions > Policies**.
3. On the Policies page, click **Create Policy**.
4. On the Create Custom Policy page, specify the **Policy Name** and the **Note**.
5. Select a **Configuration Mode**.
   - **Visualized**
     If you select the Visualized configuration mode, click **Add Statement** and configure the following parameters.

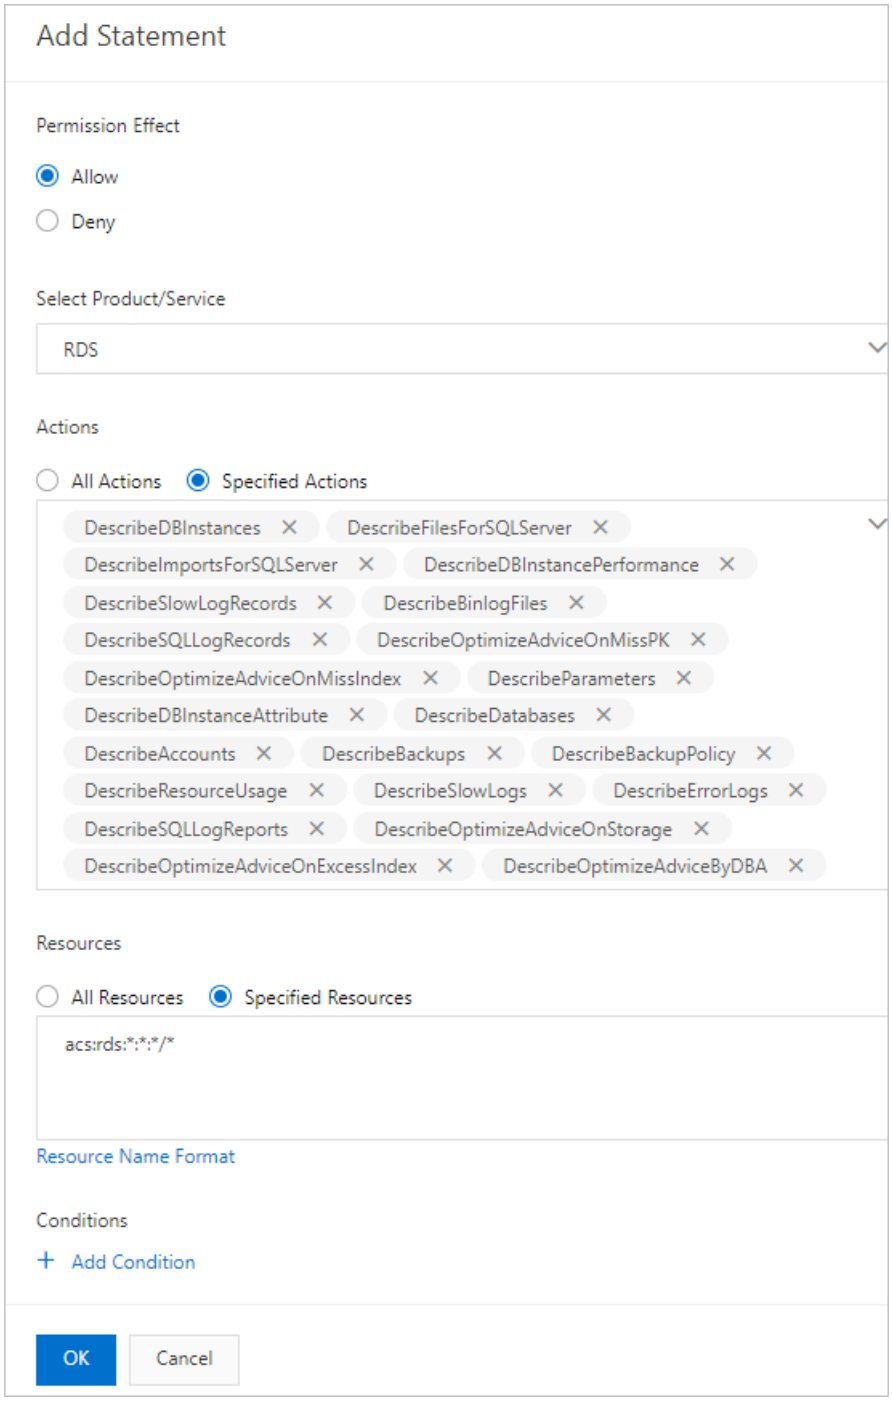

Parameter	Description
Permission Effect	Specify whether to grant the permissions on an Alibaba Cloud service to the RAM user. Valid values: Allow and Deny. In this example, select Allow.
Select Product/Service	Select the Alibaba Cloud service on which you want to grant permissions to the RAM user. In this example, select **RDS**.

Parameter	Description
Actions	Select the API operations on which you want to grant permissions to the RAM user. Valid values: **All Actions** and **Specified Actions**. If you select **Specified Actions**, you must also select the required API operations from the drop-down list that appears. In this example, select all of the API operations whose names start with **Describe**.
Resources	Select the resources on which you want to grant permissions to the RAM user. Valid values: **All Resources** and **Specified Resources**. If you select **Specified Resources**, you must also enter the names of the required resources in the following format: **acs:<service-name>:<region>:<account-id>:<relative-id>**. In this example, enter **acs:rds:*:*:*/*** . This specifies to grant the RAM user the permissions on all of the RDS instances that are created under your Alibaba Cloud account.
Conditions	Specify the limits on the permissions that you want to grant to the RAM user. For example, you can limit the source IP addresses from which the RAM user can log on.

Note If you select **Specified Actions**, we recommend that you select the **DescribeDBInstances** API operation. If you do not select this API operation, you cannot view the RDS instances in a specified region.

- **Script**
  If you select the Script configuration mode, enter the following code snippet in the edit box that appears:

```
{  "Version": "1",
   "Statement": [
    {   "Effect": "Allow",
        "Action": [
          "rds:Describe*"   ],
        "Resource": [
          "acs:rds:*:*:*/*"
          ],
        "Condition": {}
     }
  ]
}
```

> **Note** The Script configuration mode is more efficient than the **Visualized** configuration mode. For example, in Script configuration mode, you can enter `Describe*` in the edit box to specify all of the API operations whose names start with Describe. However, in Visualized configuration mode, you can select only one API operation whose name starts with Describe at a time.

6. Click **OK**.
7. In the left-side navigation pane, choose **Identities > Users**.
8. Find the RAM user, and click **Add Permissions** in the **Actions** column.
9. In the **Select Policy** section, click **Custom Policy**, find the permission policy that you created, and then click **OK**.

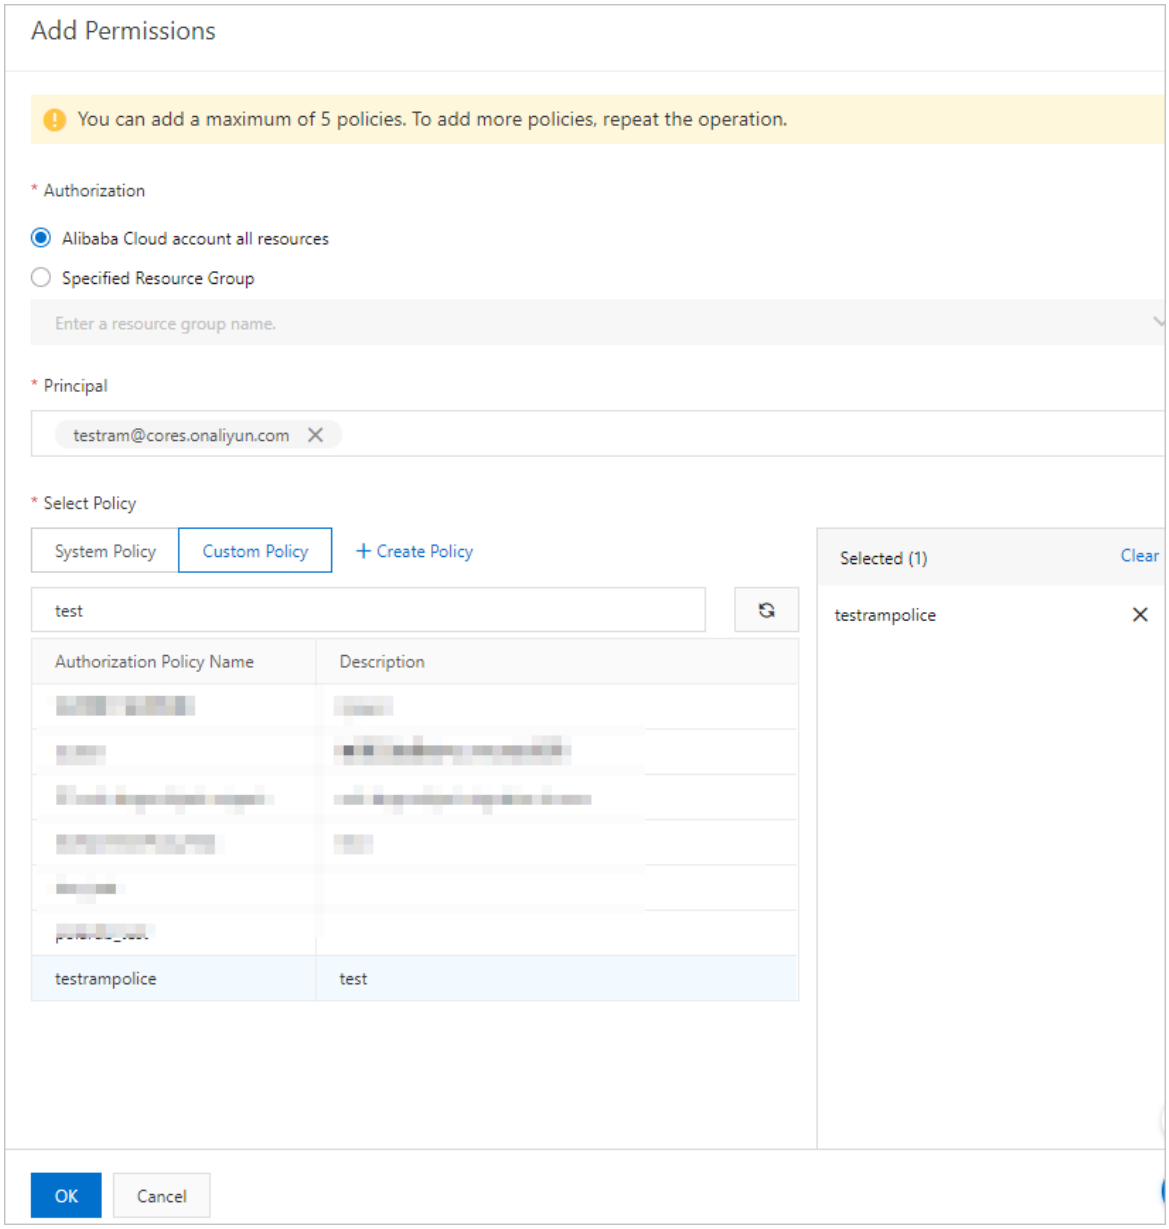

After the preceding steps are complete, you can log on to the ApsaraDB for RDS console and view RDS instances by using the credentials of the RAM user. You can also grant other permissions to a RAM user under your Alibaba Cloud account based on your business requirements.

- For more information about the API of ApsaraDB for RDS, see API overview.
- For more information about the elements of a permission policy, see Policy elements.

5.Use DTS to perform vertical splitting on a database

If an RDS instance is overloaded, you can vertically split a database or table from this instance into a separate instance. This topic describes how to use the dual-write solution and simple splitting solution to achieve vertical splitting.

Context

Assume that Database A and Database B are deployed on an RDS instance. The instance is facing bottlenecks because of business growth. If you want to reduce the load on the instance, you can vertically split Database B into a separate instance. For more information, see, Dual-write solution and Simple splitting solution.

Precautions

- You must create a separate instance as the destination instance. The database account of the destination instance must have the same permissions as the database account of the source instance. For more information about the supported destination instance types, see Overview of data migration scenarios.

  > **Note** We recommend that you create and authorize a database account for the source and destination instances. This allows you to distinguish session information and improve data security.

- You must add the connection string of the destination instance to the application.
- If you use the simple splitting solution, you must pause your business and stop writing data to the database for a short period of time. To minimize the impact on your business when you change and release a program or switch instances, we recommend that you split the database during off-peak hours.

Dual-write solution

> **Note** You can implement phased switchover and minimize the impact on the business. However, you must transform the application so that the application data can be written to Database B on the source and destination instances.

1. Configure a data migration task for the source and destination instances. Select Database B as the object to be migrated. For more information, see Overview of data migration scenarios.

   > **Note** When you configure the data migration task, you must select **Schema Migration**, **Full Data Migration**, and **Incremental Data Migration** as the migration types.

2. Wait until the task progress bar shows "Incremental Data Migration" and "The migration task is not delayed" or a delay time of less than 5 seconds.

   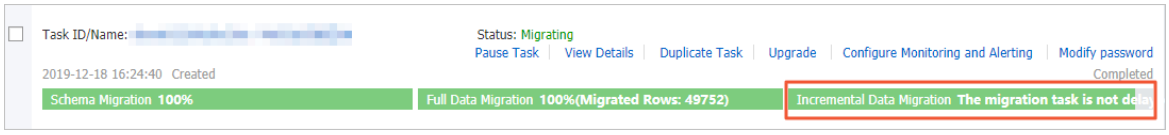

3. Check whether the data in Database B is consistent between the source and destination instances.

4. If the data is consistent, stop the data migration task. For more information, see Stop a data migration task.

   > **Warning** The database accounts that are used for data migration have the read and write permissions. After data is migrated, you must delete the accounts or revoke the write permission to ensure security.

5. Write the application data to Database B on the source and destination instances at the same time.
6. Log on to Database B on the source and destination instances and execute one of the following statements to view the session information. Select the statement based on the database type. Make sure that write operations are being performed in a new session.

   > **Note** The process or session information returned by the preceding statements includes the processes or sessions between DTS and Database B on the source and destination instances.

   MySQL | SQL Server | Oracle | PostgreSQL | Redis | MongoDB

7. Write the application data to Database B on the source and destination instances at the same time. Make sure that the source and destination instances keep running for a business period, for example, seven days.
8. After you test all features related to your business and make sure that no issues are detected, shut down Database B on the source instance.

Simple splitting solution

> **Note** If you use this solution, you do not need to edit the code of your application. However, a rollback failure may occur.

1. Configure a data migration task for the source and destination instances. Select Database B as the object to be migrated. For more information, see Overview of data migration scenarios.
2. Wait until the task progress bar shows "Incremental Data Migration" and "The migration task is not delayed" or a delay time of less than 5 seconds.

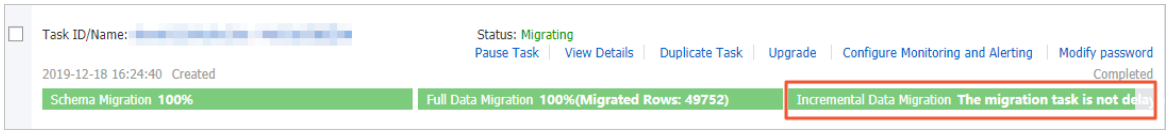

   > **Note** If you do not select **Incremental Data Migration** when you configure the data migration task, the task progress bar does not show "Incremental Data Migration". After data is migrated, the migration task automatically stops. In this case, you must pause your business and stop writing data to the source database before running the data migration task. Skip to Step 6 and proceed.

3. Check whether the data in Database B is consistent between the source and destination instances.
4. If the data is consistent, stop the data migration task. For more information, see Stop a data migration task.

> **Warning** The database accounts that are used for data migration have the read and write permissions. After data is migrated, you must delete the accounts or revoke the write permission to ensure security.

5. Disconnect the application from Database B.

> **Note** If the business of Database A is affected after the application is disconnected from Database B, you must disconnect the application from the source instance.

6. Log on to Database B on the source instance and execute one of the following statements to view the session information. Select the statement based on the database type. Make sure that no write operations are being performed in a new session.

> **Note** The process or session information returned by the preceding statements includes the process or session between DTS and Database B on the source instance.

MySQL | SQL Server | Oracle | PostgreSQL | Redis | MongoDB

7. Create and start a data migration task in the opposite direction. The task migrates incremental data generated in Database B on the destination instance to Database B on the source instance. The data migration task created in this step provides a rollback solution. If an error occurs in the destination database, you can switch workloads to the source database.

> **Warning** When you configure a data migration task in the opposite direction, you must select only "Incremental Data Migration" in the "Configure Migration Types and Objects" step. Then, you must select the database or table to be migrated back to the source database.

8. Make sure that the application is disconnected from Database B on the source instance. Verify that data is consistent between Databases B on the source instance and Databases B on the destination instance. Then, switch the database services to Databases B on the destination instance and resume your business.
9. Log on to Database B on the destination instance and execute one of the following statements to view the session information. Select the statement based on the database type. Make sure that write operations are being performed in a new session.

> **Note** The process or session information returned by the preceding statements includes the process or session between DTS and Database B on the source instance.

MySQL | SQL Server | Oracle | PostgreSQL | Redis | MongoDB

```
show processlist;
```

After you switch workloads to the destination database, make sure that the destination database keeps running for a business period, for example, seven days.

11. Test all features related to your business and make sure that no issues are detected. Then, shut down Database B on the source instance and stop the data migration task in the opposite direction. For more information, see Stop a data migration task.

Warning The database accounts that are used for data migration have the read and write permissions. After data is migrated, you must delete the accounts or revoke the write permission to ensure security.