Alibaba Cloud ApsaraDB for RDS

FAQ

Issue: 20200506

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 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 contact Alibaba Cloud directly if you discover any errors in this document.

Document conventions

Style	Description	Example
	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.
	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.
	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.
	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]`

Style	Description	Example
{} or {a|b}	This format is used for a required value, where only one item can be selected.	switch {active|stand}

Contents

1 Purchases and Payments

1.1 Why does RDS for MySQL not support the MyISAM engine?

The following lists the major reasons why RDS for MySQL does not support the MyISAM engine:

- MyISAM has defects in data integrity protection, and these defects may cause corruption or even loss of database data. Additionally, many of these defects are design issues and cannot be fixed without compromising compatibility.
- Most data corruption issues of MyISAM can only be manually fixed, and therefore MyISAM cannot be used for product services.
- For RDS storage, MyISAM is not the best solution for I/O operations. Therefore, MyISAM does not necessarily surpass InnoDB in terms of performance.
- It is easy to migrate from MyISAM to InnoDB because most applications simply need to modify the table creation code.
- MyISAM is developing towards InnoDB. MySQL 5.7 can be completely different from MyISAM and the system's data control is also switched to InnoDB.

1.2 How do I change an SSD to a local SSD?

This topic describes how to change an SSD to a local SSD for an RDS instance. You can use the same method to change a local SSD to an SSD.

To change the storage class (local SSD, SSD, or ESSD) of an RDS instance, you must use DTS to migrate data from the source RDS instance to another new RDS instance.

Prerequisites

- The DB engine is one of the following:
  - RDS for MySQL
  - RDS for SQL Server
  - RDS for PostgreSQL
- You have created an RDS instance with the storage class you want. For more information, see #unique_6.

- The storage capacity of the destination RDS instance cannot be lower than that of the source RDS instance.
- The destination RDS instance is located in the same region as the source RDS instance.
- The DB engine, version, and edition of the destination RDS instance are the same as those of the source RDS instance.
- If you select the incremental data migration type, the **binlog_row_image** parameter is set to full for the source RDS instance.

Precautions

- The instance information changes after the migration, therefore you must modify the instance information on your application properly to guarantee service continuity.
- Do not perform DDL operations during the migration.
- Event migration is not supported in schema migration.
- If the object name mapping function is used for an object, the migration of objects relying on the object may fail.

Procedure

1. Log on to the DTS console.
2. In the left-side navigation pane, click **Data Migration**. Then in the upper-right corner, click **Create Migration Task**.

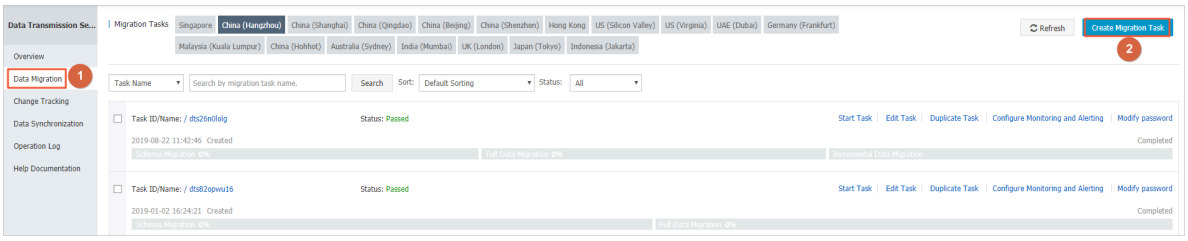

3. Set the following parameters.

Parameter		Description
Task Name		By default, DTS automatically generates a name for each task. You can change the name to indicate the specific services for easy identification of the task.
Source Database	**Instance Type**	Select **RDS Instance**.
	Instance Region	Select the region where the source RDS instance is located.
	RDS Instance ID	Select the ID of the source RDS instance.

Parameter		Description
	Database Account	Enter the username of the database account you use to connect the source RDS instance. For example, the account can be a premier account, or a standard account that has the read and write permissions for all databases.
	Database Password	Enter the password of the database account you use to connect the source RDS instance.
	Encryption	In typical cases, select **Non-encrypted**. If the source RDS instance supports #unique_7 and has SSL encryption enabled, select **SSL-encrypted**.
Destination Database	**Device Spec**	Select **RDS Instance**.
	Instance Region	Select the region where the source RDS instance is located.
	RDS Instance ID	Select the ID of the destination RDS instance.
	Database Account	Enter the username of the database account you use to connect the destination RDS instance. For example, the account can be a premier account, or a standard account that has the read and write permissions for all databases.
	Database Password	Enter the password of the database account you use to connect the source RDS instance.
	Encryption	In typical cases, select **Non-encrypted**. If the source RDS instance supports #unique_7 and has SSL encryption enabled, select **SSL-encrypted**.

Note:

The values of the **Instance Type** and **RDS Instance ID** parameters determine which of the other parameters are displayed.

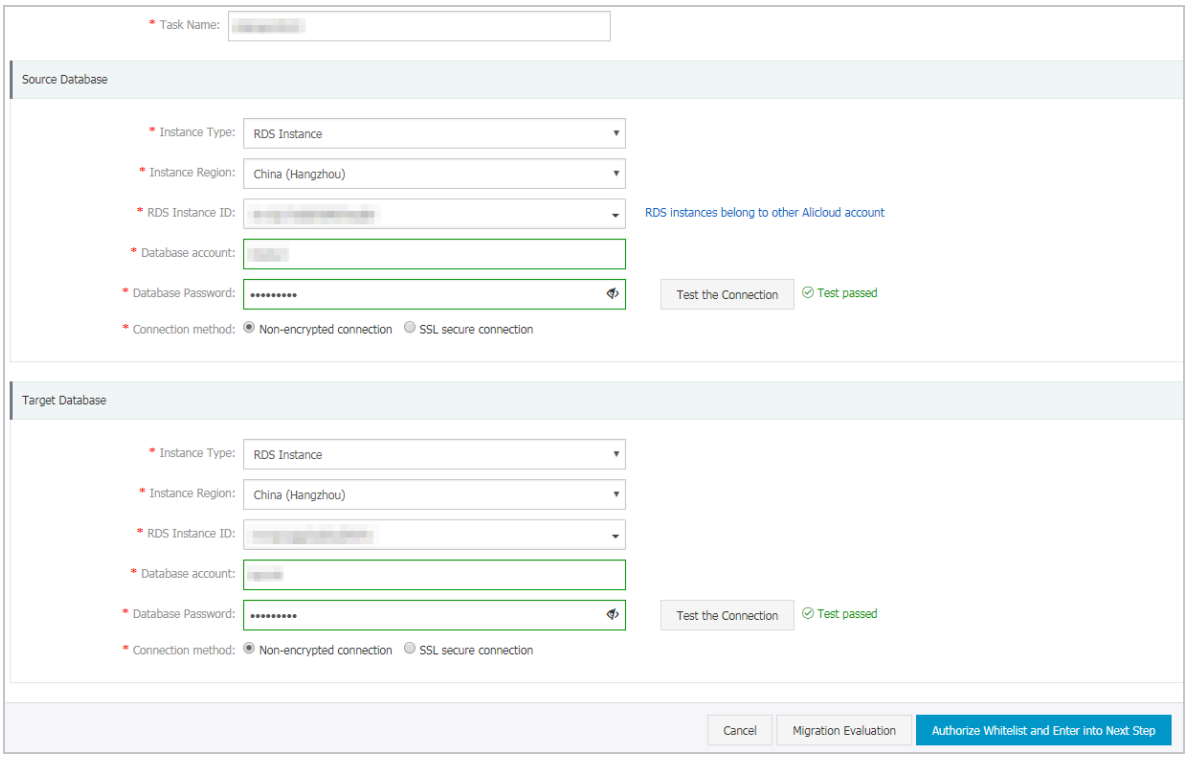

4. Click **Test Connectivity** in the **Source Database** and **Destination Database** sections separately.

   Note:

   If a message is displayed to the right of **Test Connectivity**, the RDS instance can be connected. Otherwise, rectify faults based on the error message.

5. In the lower-right corner, click **Set Whitelist and Next**.

6. Select a migration type. In the **Available** section, select the objects you want to migrate, and click **>** to add them to the **Selected** section. Then, click **Precheck**.

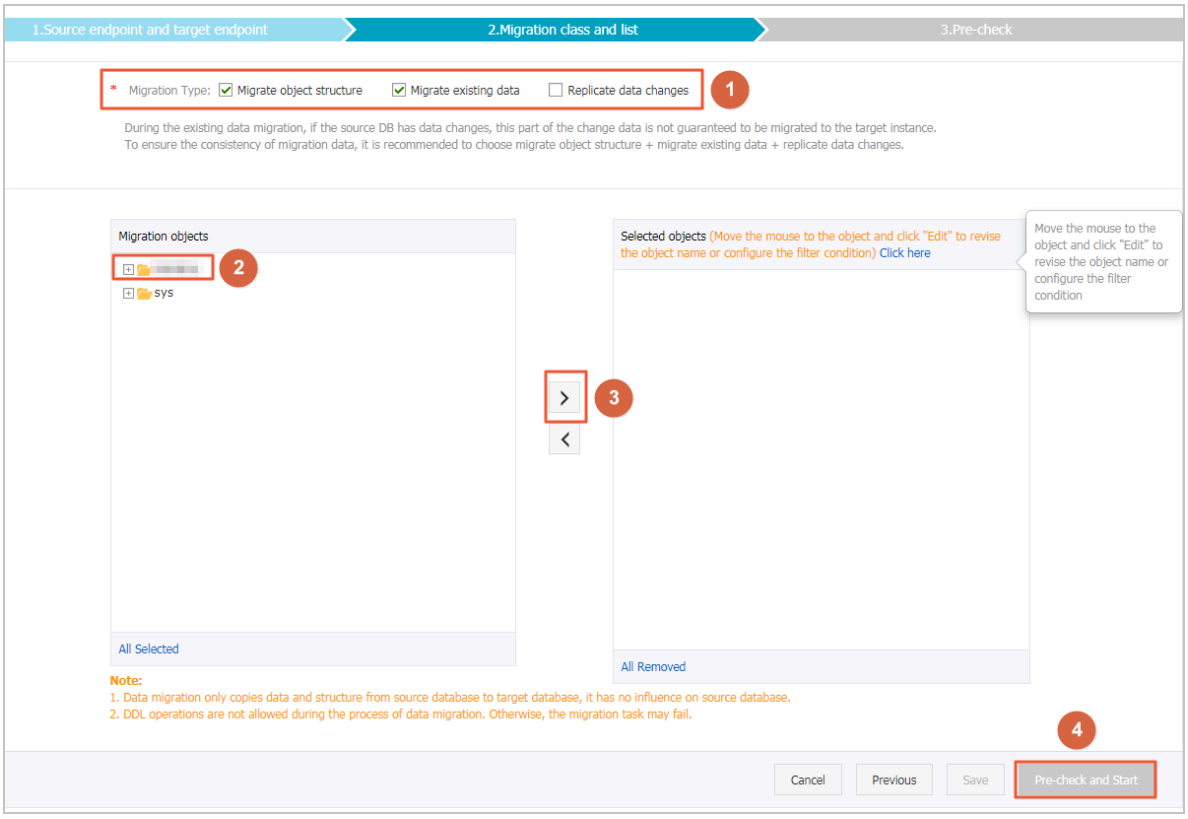

Note:

If you want to change the name of an object to be migrated, then find the object in the **Selected** section and click **Edit** to the right of the object.

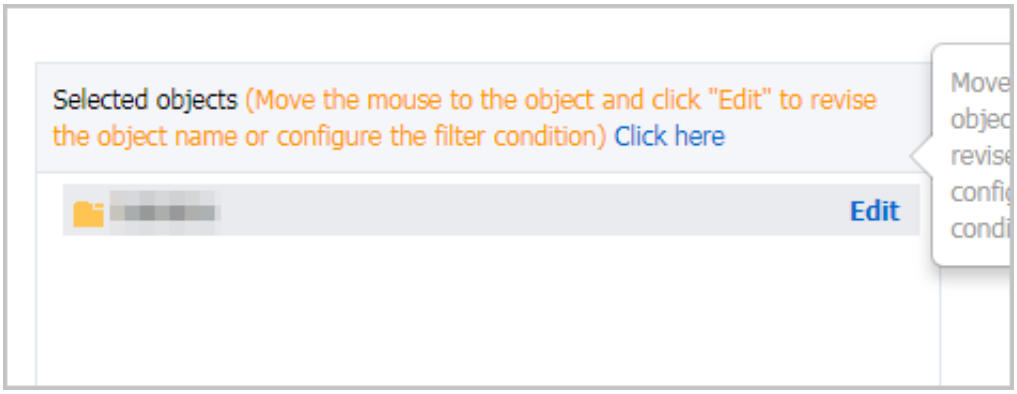

7. Optional. If the migration task fails the precheck, perform this step. If the migration task passes the precheck, go to Step 10.

   The system displays the precheck results, as shown in the following figure.

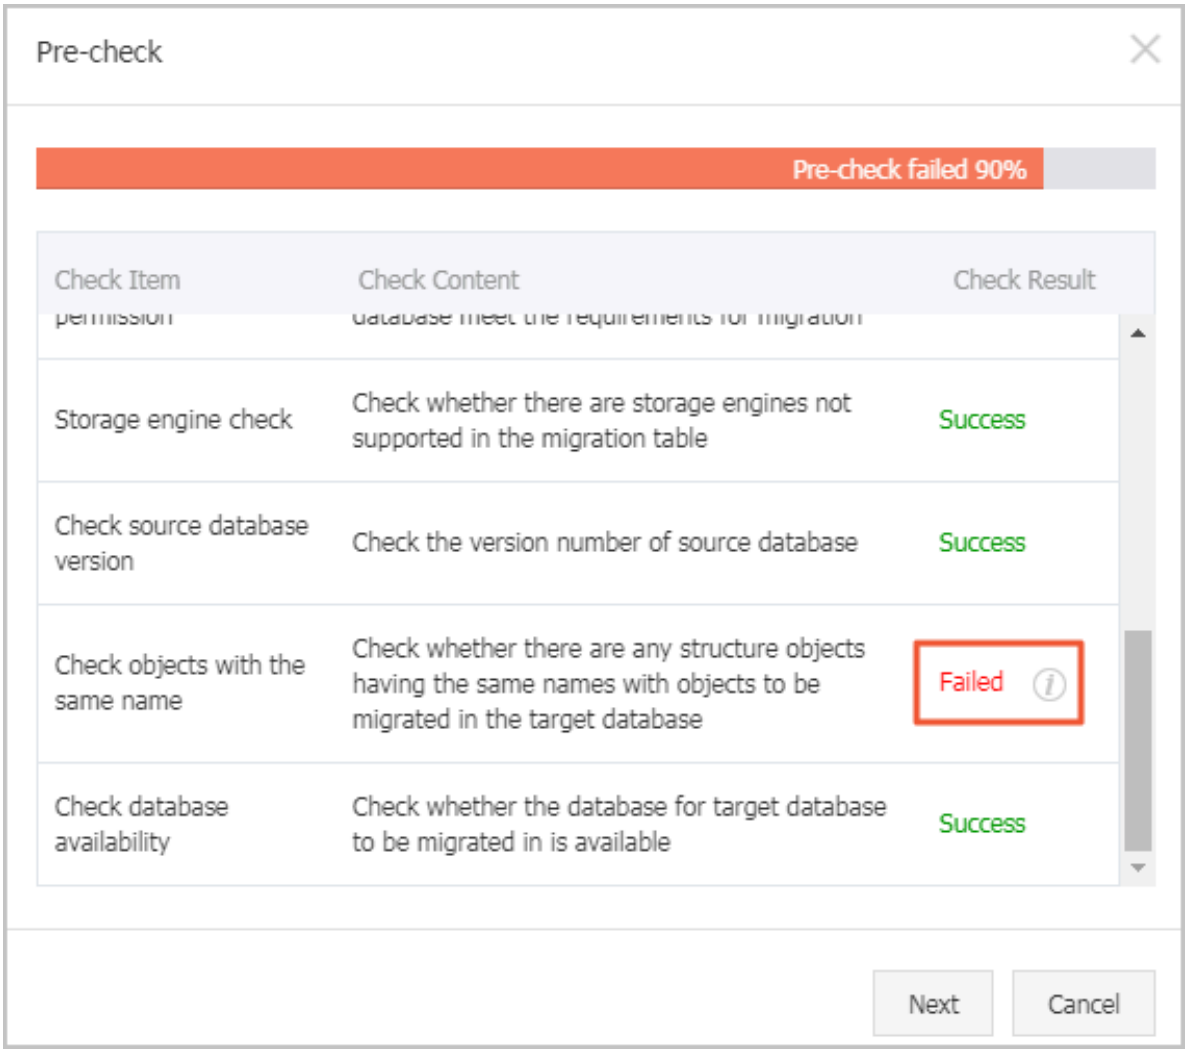

8. Find each check item whose Check Result**Check Result** is **Failed** and click ⓘ to view the failure details. Then locate the fault based on the failure details.

9. After all failures are located, navigate to the page that displays the migration task list, and start the migration task you have created.

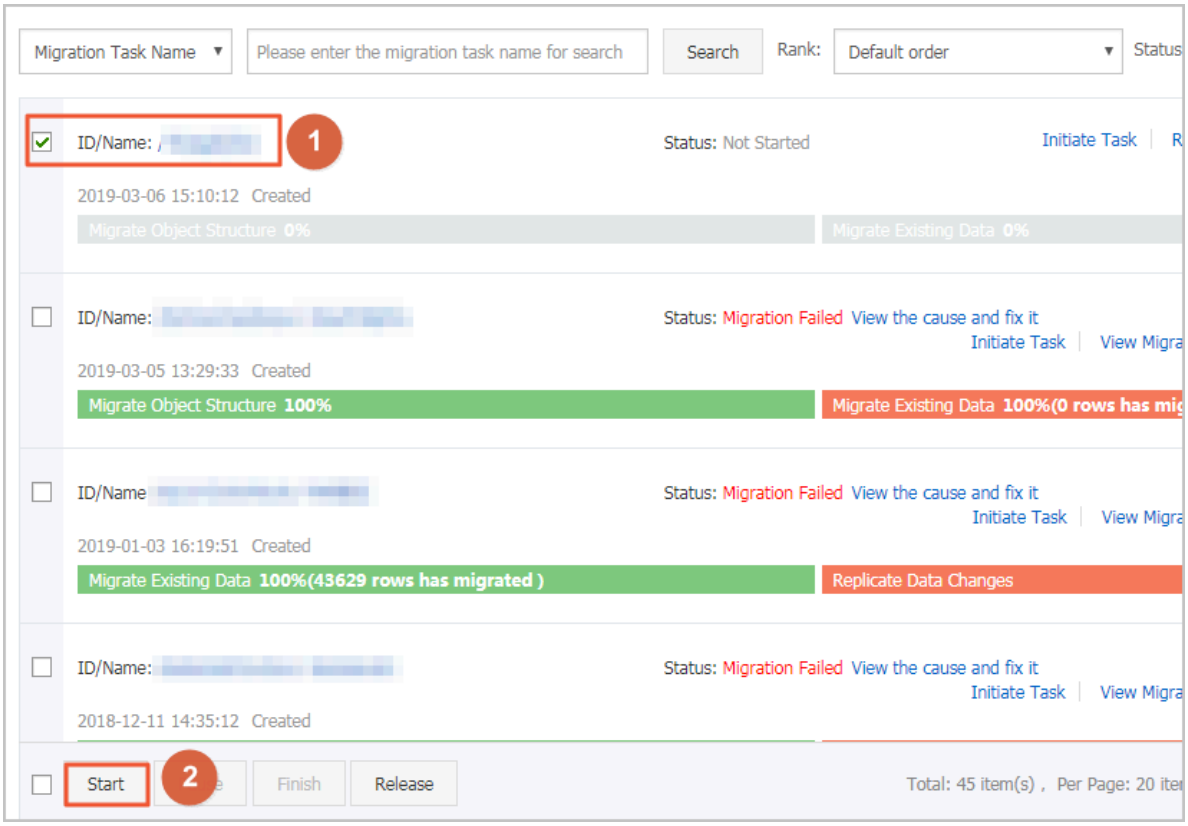

10. When the migration task passes the precheck, click **Next**.

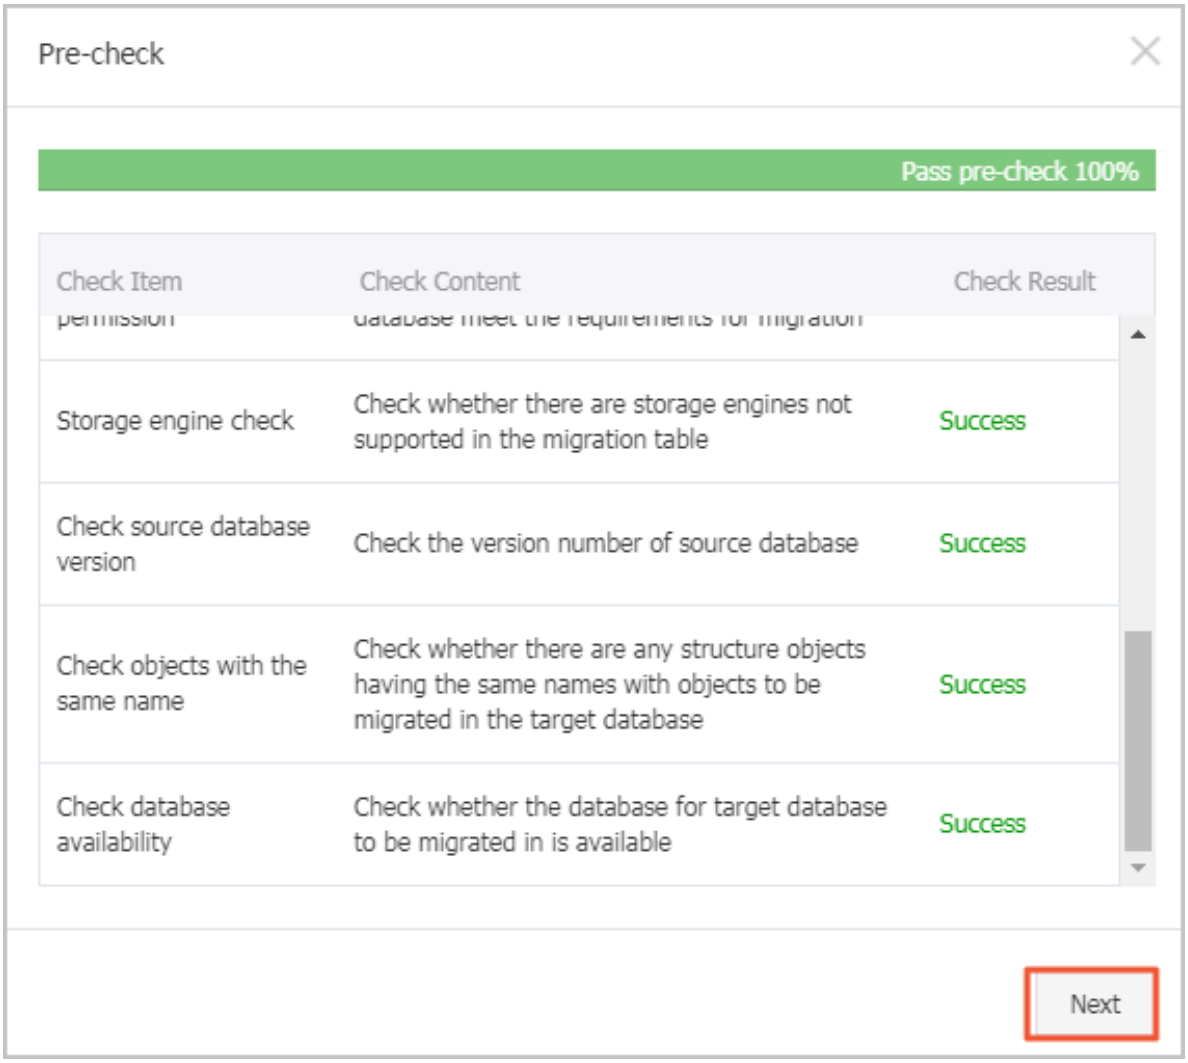

11. Confirm the settings, select **Data Transmission Service (Pay-As-You-Go) Service Terms**, and click **Buy and Start**.

What to do next

Wait until the migration is complete. After the migration is complete, you must modify the instance information on your application and use the endpoint of the new RDS instance to establish a connection.

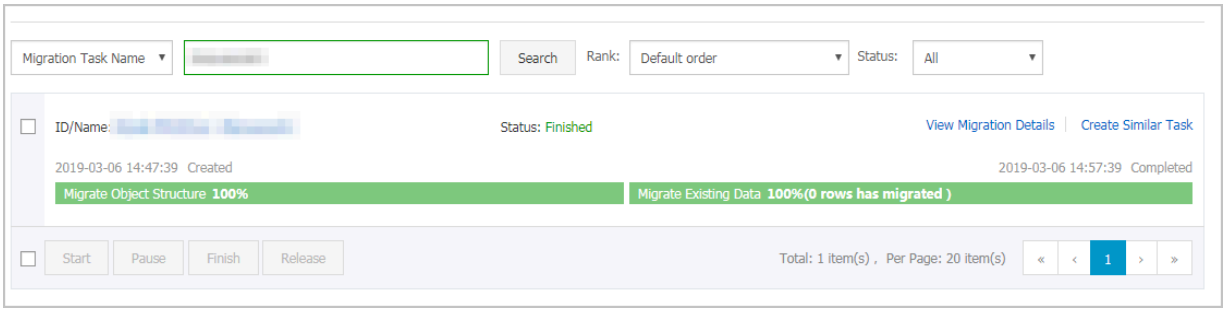

2 Connections and Networks

2.1 How do I locate the public IP address of my computer that needs to connect to RDS for MySQL or MariaDB TX?

Problem

- You have added the public IP address of your computer to the IP address whitelist of the RDS instance. However, your computer cannot access the instance while other devices can.
- You have added the public IP address of your computer to the IP address whitelist of the RDS instance, but your computer cannot access the instance unless you set the IP address whitelist to 0.0.0.0/0 or your company's address range.

If either of the preceding problems occurs, the public IP address you add to the whitelist of the RDS instance may be incorrect. You need to find the real public IP address of your computer.

> **Note:**
> This topic applies only when you access the RDS instance from a device other than an ECS instance. If you access the RDS instance from an ECS instance, you can find the public and private IP addresses of the ECS instance on the ECS console.

Precautions

If the public IP address of your computer is not fixed and your RDS instance is used in a production environment, we recommend that you use a private connection instead or configure an appropriate IP address range in the whitelist. This is to ensure that the connection remains available even if the public IP address of your computer changes.

Procedure

1. Add your company's public IP address range or 0.0.0.0/0 to the IP address whitelist of the RDS instance. For more information, see #unique_10.

   > **Note:**

0.0.0.0/0 indicates that all devices are allowed to access the RDS instance. Use the address with caution. If you add 0.0.0.0/0 to the whitelist, we recommend that you remote it from the whitelist immediately once you no longer need it.

2. Connect your computer to the RDS instance by using a client or the command line interface (CLI).

```
mysql -h<RDS Connection address> -u<Username> -p<Password> -P3306
```

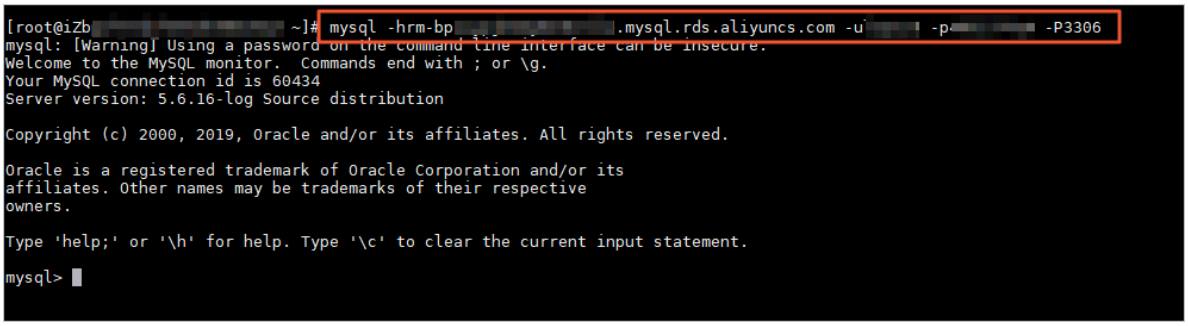

3. Check the process information.

```
show processlist
```

As shown in the following figure, the value of **Host** for the **show processlist** record is the real public IP address of your computer.

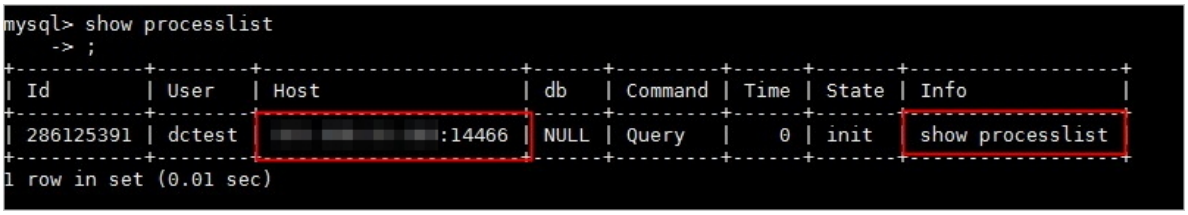

4. Remove 0.0.0.0/0 from the whitelist and add the real public IP address of your computer to the whitelist.

2.2 Navicat connects to databases such as apsaradb RDS for MySQL

This topic describes how to use Navicat to connect to RDS databases, such as MySQL, SQL Server, PostgreSQL, PPAS, and MariaDB.

Note:

Alibaba Cloud has officially launched Data Management Service (DMS), which provides graphical Data Management functions. It can be used to manage relational databases and NoSQL databases, and supports Data Management, structure Management, user

authorization, provides features such as security audit, data trend analysis, data tracing, BI charts, and performance and optimization.

We recommend that you use DMS to manage databases. For more information about how to use DMS to connect to a database, see #unique_12.

Prerequisites

You have completed the following operations:

- Create a cluster
- #unique_10
- Create a RAM user

Procedure

1. Start Navicat.
2. In the upper-left corner, **connection** to select the type of the database to be connected.

Note:

- If no **alibaba Cloud** menu (Navicat version is too low), you can directly select **MySQL**, **PostgreSQL** and other databases.
- When connecting to a PPAS instance, Select **PostgreSQL**.

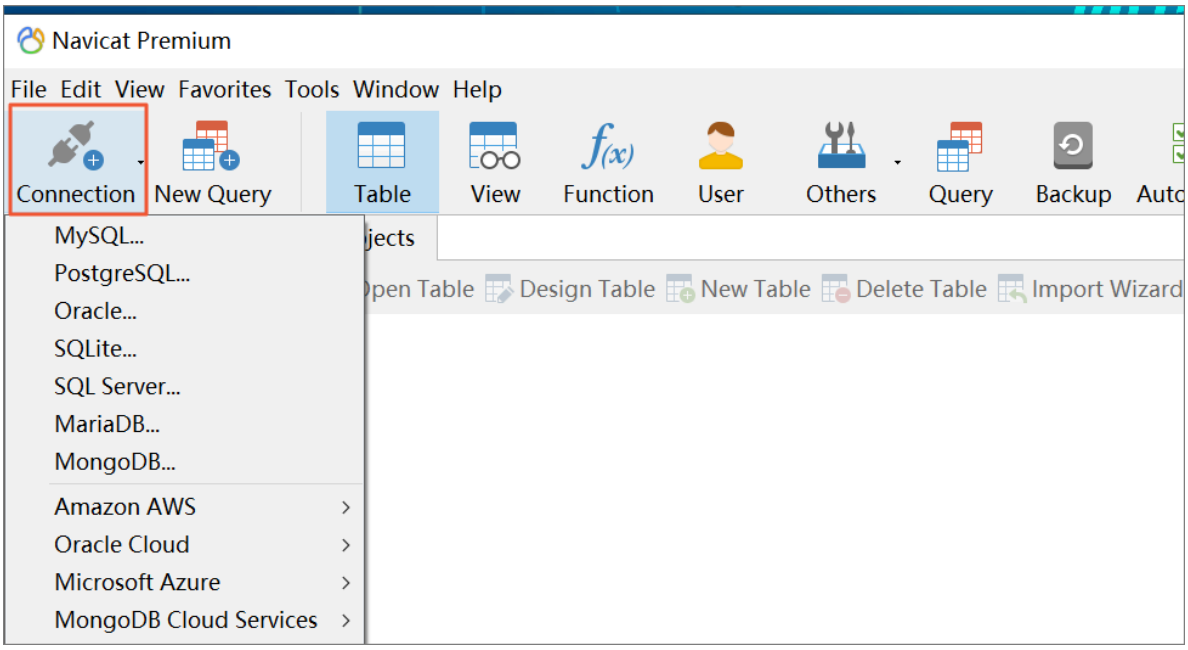

3. Enter information about the RDS instance you want to connect. The parameters are described as follows.

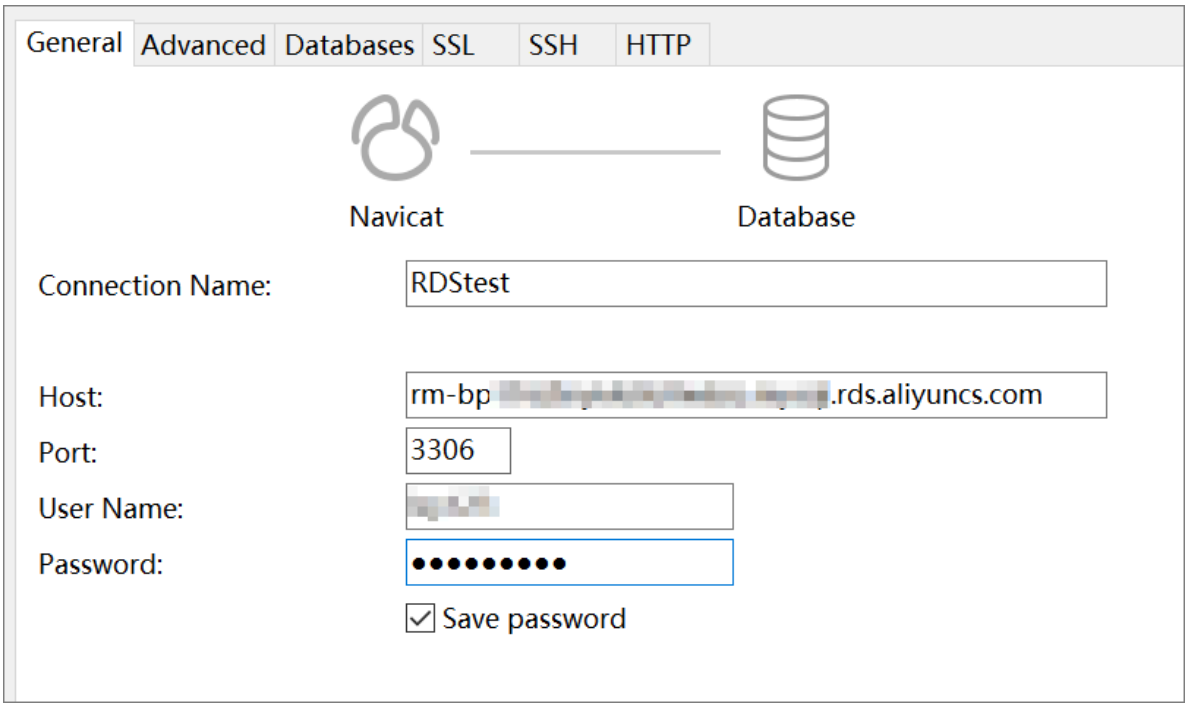

Parameter or button	DO NOT TRANSLATE
Connection Name	The custom name of the connection.
Host	Enter the intranet address or internet address of the RDS instance. Instructions for using the intranet address and Internet address are as follows: • If your client is deployed in an ECS instance, and the instance is in the same region and has the same network type as the destination RDS instance, you can use the internal IP address. For example, ECS and RDS instances are both in the VPC located in China (Hangzhou). You can use the internal IP address provided to create secure connection. • Use the public IP address for other situations. For more information about how to view the internal network endpoint and public network endpoint, see #unique_15.
Port	If you connect to the instance over the internal network, enter the internal port of the instance. If you connect to the instance over the public network, enter the public port of the instance.
UID	Enter the name of the account used to access the RDS instance.

Parameter or button	DO NOT TRANSLATE
Password	**Username** the password.

4. Click **OK**.

> **Note:**
> If you cannot connect to the instance, see #unique_16.

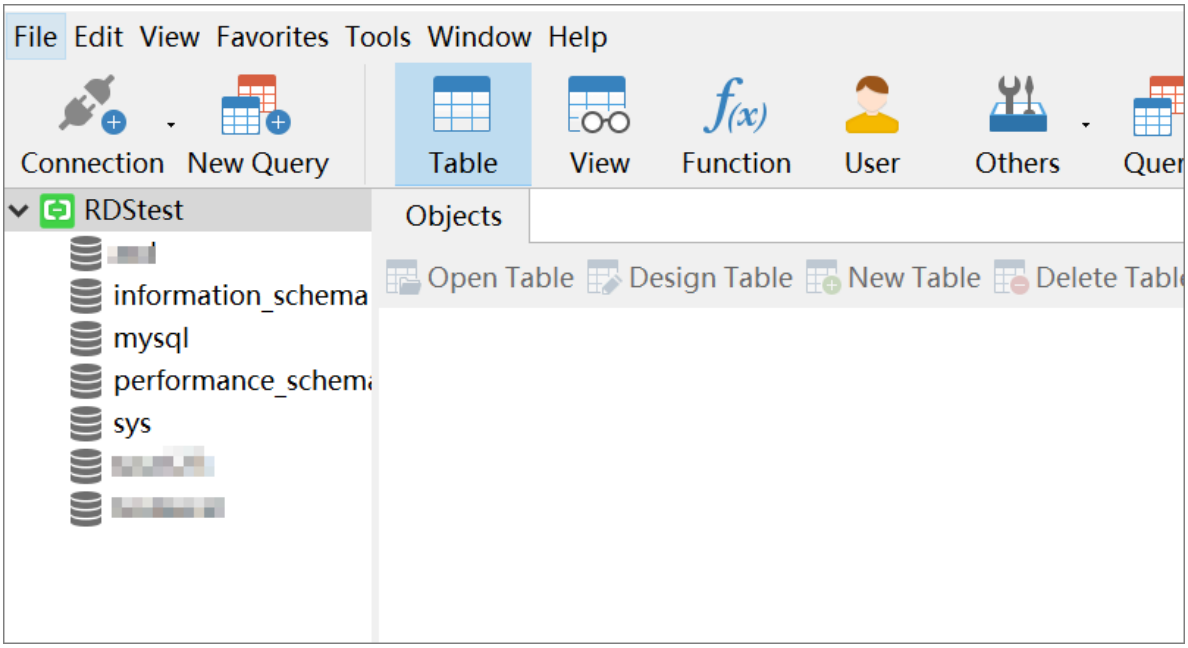

2.3 How do I access an RDS instance from an ECS instance over the intranet?

If your business is built on an ECS instance and you need to access an RDS instance from the ECS instance, we recommend that you establish the connection over the intranet. This increases the transmission rate and ensures security.

For more information, see How do I connect to an ApsaraDB for RDS instance?.

2.4 How do I connect to an ApsaraDB for RDS instance?

This topic describes how to connect to an ApsaraDB for RDS instance over the Internet or an internal network. We recommend that you establish a connection over an internal network to ensure data security and transmission efficiency.

Connect to an RDS instance over the Internet

Log on to the ApsaraDB for RDS console, find the RDS instance, and navigate to the **Database Connection** page to apply for a public endpoint. By default, an RDS instance is not provided with a public endpoint.

Note:

- For security purposes, exercise caution when you connect to an RDS instance by using its public endpoint.
- For faster transmission and higher data security, we recommend that you migrate your application to an ECS instance located in the same region and with the same network type as your RDS instance and then connect to your RDS instance by using its internal endpoint.

After you obtain a public endpoint, you can use it to connect to your RDS instance. For more information, see References.

Connect to an RDS instance over an internal network

Log on to the ApsaraDB for RDS console, find the RDS instance, and navigate to the **Database Connection** page to obtain the internal endpoint of the RDS instance.

Requirements

In normal cases, you can only connect to your RDS instance over an internal network from an ECS instance. To connect to your RDS instance from an on-premises data center, you must establish a physical connection between them. For more information, see Connect an on-premises data center to a VPC through a physical connection.

When you connect to your RDS instance over an internal network from an ECS instance, make sure that the following requirements are met:

- The ECS and RDS instances are created by the same Alibaba Cloud account.
- The ECS and RDS instances reside in the same region.
- The ECS and RDS instances have the same network type.

- If the ECS and RDS instances both have the VPC network type, they must reside in the same VPC.
- The private IP address of the ECS instance is added to an IP address whitelist of the RDS instance. For more information, see Configure a whitelist.

If all the preceding requirements are met, you can connect to the RDS instance by using its internal endpoint from the ECS instance. For more information, see References.

FAQ

- How do I disable the connection established over the Internet to an RDS instance?

  Make sure that the whitelists of the RDS instance only contain private IP addresses. Alternatively, you can release the public endpoint of the RDS instance on the **Database Connection** page.
- Why am I unable to change the network type of an RDS instance from VPC to Classic Network?

  The RDS instance only supports the VPC network type. For more information, see #unique_22.
- Why is my RDS instance disconnected from an ECS instance even though the public IP address of the ECS instance is added to an IP address whitelist of my RDS instance?

  The public IP address of the ECS instance may have changed. In such cases, you must add the new public IP address of the ECS instance to an IP address whitelist of your RDS instance.

References

- Connect to an RDS for MySQL instance
- Connect to an RDS for SQL Server instance
- Connect to an RDS for PostgreSQL instance
- Connect to an RDS for PPAS instance
- Connect to an RDS for MariaDB instance

2.5 How do I locate the IP address connected to an RDS for SQL Server instance?

Obtain the IP address of your computer connected to an RDS instance

Problem description

The public IP address of your computer dynamically changes, therefore the IP address you obtain by using a local IP address query tool may be incorrect. As a result, RDS reports connection errors even after you add the obtained public IP address to the IP address whitelist of the RDS instance. You can access the RDS instance only after you obtain the correct IP address of your computer.

Precautions

If the public IP address of your computer changes and the established connection to the RDS instance is used in a production environment, we recommend that you use a private network connection instead or add an appropriate CIDR block to the IP address whitelist of the RDS instance. This helps to guarantee a stable connection despite changes to the public IP address of your computer.

Procedure

1. Add the IP address **0.0.0.0/0** to the IP address whitelist of the RDS instance. For more information, see #unique_29.

    **Note:**
   The IP address **0.0.0.0/0** indicates that all IP addresses are allowed to access the RDS instance.

2. Use a client to connect your computer to the RDS instance. For more information, see #unique_24.
3. Run the following commands to query the IP address of your computer:

   ```
   SELECT  CONNECTIONPROPERTY('PROTOCOL_TYPE') AS PROTOCOL_TYPE,
         CONNECTIONPROPERTY('CLIENT_NET_ADDRESS') AS CLIENT_NET_ADDRESS
   ```

4. Delete the IP address **0.0.0.0/0** that you added to the IP address whitelist in Step **1**, and add the real outbound IP address of your computer to the IP address whitelist.

Obtain all IP addresses connected to an RDS instance

Problem description

You want to obtain all IP addresses that are connected to the RDS instance, or you want to locate security issues such as link leakage.

Procedure

1. Add the IP address **0.0.0.0/0** to the IP address whitelist of the RDS instance. For more information, see#unique_29.

2. Use a client to connect your computer to the RDS instance.
3. Run the following commands to query all IP addresses that are connected to the RDS instance.

```
SELECT
SP.SPID,
SP.LOGINAME,
SP.LOGIN_TIME,
SP.HOSTNAME,
SP.PROGRAM_NAME,
DC.CLIENT_TCP_PORT,
DC.CLIENT_NET_ADDRESS
FROM SYS.SYSPROCESSES AS SP
INNER JOIN SYS.DM_EXEC_CONNECTIONS AS DC
ON SP.SPID = DC.SESSION_ID
WHERE SP.SPID > 50
AND DC.AUTH_SCHEME='SQL'
```

4. Delete the IP address **0.0.0.0/0** or the CIDR block containing your company's IP address segment that you added to the IP address whitelist in Step **1**.

View the parameter settings of a connection

After you obtain all IP addresses that are connected to the RDS instance, you can run the following command to view the parameter settings of a specific connection:

```
SELECT * FROM SYS.DM_EXEC_SESSIONS WHERE SESSION_ID=<The obtained SPID>
```

3 Migration and Synchronization

4 Backup and Restoration

4.1 Differences between data migration, synchronization, and recovery

- Comparison between data restoration and migration

  Although data restoration and migration can both be used to replicate data to a new RDS instance, they differ in the following aspects:

  - The data restoration feature enables you to restore data from a data backup file, and restore data from a specific point in time by using a data backup file and a binary log file.
  - The data migration feature enables you to use Alibaba Cloud Data Transmission Service (DTS) to migrate data from an on-premises database, a user-created database on an ECS instance, or a database on a third-party cloud platform to Alibaba Cloud. This feature also enables you to migrate data from Alibaba Cloud to a user-created database.

  You can use data restoration and migration together to meet more business needs. For example, you can restore data to a new RDS instance and then use DTS to migrate the data back to its original RDS instance. This way, you do not need to update the endpoints on your application.

  Note:
  The data restoration feature provided in the ApsaraDB for RDS console can meet most business needs for data restoration and some business needs for data migration. However, it does not support incremental data migration.

- Comparison between data migration and synchronization

  - The data migration feature is used to migrate data to Alibaba Cloud. If you select the incremental data migration option when you create a migration task, the data updated in the original RDS instance is synchronized to the new RDS instance. However, this feature is different from the data synchronization feature, which

supports more functions such as online updates of synchronized objects and two-way data synchronization. After a data migration task is complete, you can release it.

- The data synchronization feature is used to synchronize data between two data sources in real time. It is suitable for diversified scenarios such as active geo-redundancy, disaster recovery, cross-border data synchronization, query load balancing, cloud BI systems, and real-time data warehousing. After you create a data synchronization task, it keeps synchronizing data between the specified data sources to ensure data consistency.

4.2 Use DMS to export data to a CSV file

This topic describes how to use DMS to export data to a CSV file.

Check source data

In DMS, check that the data is displayed normally.

View the character set of a table

Execute the following statement to create a table:

```
show create table <table name>;
```

```
mysql>show create table st
+-----------------+-------------------------------------------------------------------
| Table           | Create Table
+-----------------+-------------------------------------------------------------------
| st              | CREATE TABLE `st` (
  `id` int(11) DEFAULT NULL,
  `name` varchar(10) COLLATE utf8mb4_unicode_ci DEFAULT NULL,
  `age` int(11) DEFAULT NULL,
  `column_name` varchar(191) COLLATE utf8mb4_unicode_ci DEFAULT NULL
) ENGINE=InnoDB DEFAULT CHARSET=utf8mb4 COLLATE=utf8mb4_unicode_ci |
+-----------------+-------------------------------------------------------------------
```

Create an export task

1. In DMS, choose **Data Operation** > **Export**.
2. Choose **New** > **Export Database**.
3. Select the table to be exported, set **File Type** and **Character Set**, and click **OK**.

4.3 mysqldump options for ApsaraDB RDS for MySQL

GTID feature

The GTID feature was introduced in ApsaraDB RDS for MySQL 5.6 to add the set-gtid-purged option to the mysqldump utility bundled with MySQL 5.6.

Option	Default value	Optional value	Purpose
set-gtid-purged	AUTO	ON|OFF|AUTO	Specifies whether to include a SET @@GLOBAL.GTID_PURGED statement in the mysqldump output.

Note:

- ON: A SET @@GLOBAL.GTID_PURGED statement is included in the mysqldump output.
- OFF: A SET @@GLOBAL.GTID_PURGED statement is not included in the mysqldump output.
- AUTO: the default value. For an instance with GTID enabled, a SET @@GLOBAL.GTID_PURGED statement is included in the mysqldump output. For an instance that is not started or does not support GTID, no GTID information is generated.

When using the mysqldump utility that comes with MySQL 5.6 or later to dump data from an ApsaraDB RDS for MySQL 5.5 instance, you must set set-gtid-purged to OFF. Otherwise, one of the following errors will be returned:

```
Error: Server has GTIDs disabled.
or
mysqldump: Couldn't execute 'SELECT @@GTID_MODE': Unknown system variable '
GTID_MODE' <1193>
```

Avoid table-level lock wait

In mysqldump, the lock-tables option is enabled by default to lock all tables before dumping them and prevent DML operations on these tables.

The InnoDB and TokuDB engines supported by ApsaraDB RDS for MySQL instances support transactions. We recommend that you use the single-transaction option to dump tables, rather than setting the lock-all-tables or lock-tables option.

Option	Default value	Optional value	Purpose
lock-all-tables	FALSE	TRUE|FALSE	Specifies whether to lock all tables across all databases. This is achieved by acquiring a global read lock for the duration of the whole dump. All tables in all databases are read-only during the dump. This option automatically disables the lock-tables and single-transaction options. ApsaraDB for RDS does not support this option.
lock-tables	TRUE	TRUE|FALSE	Specifies whether to lock tables before dumping them. By default, this option is enabled. You can specify the skip -lock-tables option to disable this option.
single-transaction	FALSE	TRUE|FALSE	Specifies whether to start a transactio n before dumping data. This option automatically disables the lock-tables option.

Set the character set for data to be dumped

If no character set is specified, mysqldump uses the UTF-8 character set to dump data.

Option	Default value	Optional value	Purpose
default-character-set	UTF8	Character sets supported by ApsaraDB RDS for MySQL instances	Sets the default character set that mysqldump uses to dump data.

Additional options to be noted when data is dumped

Option	Default value	Optional value	Purpose
no-defaults	N/A	N/A	Reads the .mylogin.cnf option file only.
defaults-file= file_name	N/A	N/A	Reads a specified option file.

Option	Default value	Optional value	Purpose
add-drop-database	FALSE	TRUE|FALSE	Specifies whether to add a DROP DATABASE statement before each CREATE DATABASE statement.
add-drop-table	TRUE	TRUE|FALSE	Specifies whether to add a DROP TABLE statement before each CREATE TABLE statement. By default, this option is enabled. You can use the skip-add-drop-table option to disable this option.
add-locks	TRUE	TRUE|FALSE	Specifies whether to surround each table-related statement with the LOCK TABLES tab_name WRITE and UNLOCK TABLES statements. This results in faster inserts when the dump file is reloaded.
compatible=name	N/A	ansi|postgresql|oracle|mssql	Enhances compatibility with a specified database type.
compact	FALSE	TRUE|FALSE	Specifies whether to enable the skip-add-drop-table, skip-add-locks, skip-comments, skip-disable-keys, and skip-set-charset options.
databases	TRUE	TRUE|FALSE	Specifies whether to dump multiple databases. By default, mysqldump treats the first name argument on the command line as a database name and following names as table names. When this option is specified, mysqldump interprets all name arguments as database names and includes the CREATE DATABASE and USE DATABASE statements in the output before each new database.

Option	Default value	Optional value	Purpose
disable-keys	TRUE	TRUE|FALSE	Specifies whether to surround INSERT statements with `/! 40000 ALTER TABLE tab_name DISABLE KEYS /` and `/! 40000 ALTER TABLE tab_name ENABLE KEYS /` statements to speed up dump file loading. This option is effective only for non-unique indexes of MyISAM tables.
events	FALSE	TRUE|FALSE	Specifies whether to dump scheduled events from dumped databases.
extended-insert	TRUE	TRUE|FALSE	Specifies whether to combine single-row INSERT statements into a single statement that inserts multiple table rows.
hex-blob	FALSE	TRUE|FALSE	Specifies whether to export BINARY, VARBINARY, and BLOB types of data in hexadecimal notation. We recommend that you add this option when migrating data between different database versions.
ignore-table=db.tab	TRUE	TRUE|FALSE	Specifies whether to ignore a table or view. Format: database name.table name (db.tab). You can use this option multiple times to ignore multiple tables.
max-allowed-packet	24MB	24MB-1GB	Specifies the maximum size of the buffer for communication between mysqldump and an ApsaraDB RDS for MySQL instance. The default value is 24MB. The maximum value is 1GB.
no-create-db	FALSE	TRUE|FALSE	Specifies whether to exclude CREATE DATABASE statements from the output.
no-create-info	FALSE	TRUE|FALSE	Specifies whether to exclude CREATE TABLE statements from the output.
no-data	FALSE	TRUE|FALSE	Specifies whether to export an ApsaraDB RDS for MySQL database to a dump file without data.

Option	Default value	Optional value	Purpose
opt	TRUE	TRUE|FALSE	Specifies whether to enable the add-drop-table, add-locks, create-options, disable-keys, extended-insert, lock-tables, quick, and set-charset options. You can specify the skip-opt option to disable the opt option enabled by default.
dump-date	TRUE	TRUE|FALSE	Specifies whether to add a dump date to the end of the output if the comments option is specified. By default, the comments option is enabled.
routines	FALSE	TRUE|FALSE	Specifies whether to dump stored procedures and functions from dumped databases. By default, stored procedures and functions are not dumped.
result-file	TRUE	TRUE|FALSE	Specifies whether to direct output to a specified file.
set-charset	TRUE	TRUE|FALSE	Specifies whether to add SET NAMES default_character_set to the output.
triggers	TRUE	TRUE|FALSE	Specifies whether to include triggers for each dumped table in the output.

Options that are not supported by ApsaraDB RDS for MySQL instances

Option	Default value	Optional value	Purpose
all-databases	FALSE	Character sets supported by ApsaraDB RDS for MySQL instances	Specifies whether to dump all databases, including the mysql database.
flush-logs	FALSE	TRUE|FALSE	Specifies whether to execute the `flush logs;` statement in an ApsaraDB RDS for MySQL instance before starting the dump.

Option	Default value	Optional value	Purpose
flush-privileges	FALSE	TRUE|FALSE	Specifies whether to add a `flush privileges;` statement to the dump output after dumping the mysql database.
lock-all-tables	FALSE	TRUE|FALSE	Specifies whether to lock all tables across all databases. This is achieved by acquiring a global read lock for the duration of the whole dump. All tables in all databases are read-only during the dump. This option automatically disables the lock-tables and single-transaction options.
tab=dir_name	N/A	N/A	Produces tab-separated text-format data files in a specified directory. For each dumped table, mysqldump creates a tbl_name.sql file that contains the CREATE TABLE statement for creating the table, and the server writes a tbl_name.txt file that contains table data.

Unsupported options and reasons

- all-databases: Common ApsaraDB RDS for MySQL users cannot export all database tables because they do not have operation permissions on some of the tables in the mysql database.

  Error message

  ```
  mysqldump: Couldn't execute 'show create table slow_log': SHOW command denied to user 'xxx'@'xx.xx.xx.xx' for table 'slow_log' (1142)
  ```

- flush-logs: Common ApsaraDB RDS for MySQL users cannot execute the `flush logs;` statement because they do not have the RELOAD privilege.

  Error message

  ```
  mysqldump: Couldn't execute 'FLUSH TABLES': Access denied; you need (at least one of) the RELOAD privilege(s) for this operation (1227)
  ```

- flush-privileges: This option is not needed because ApsaraDB RDS for MySQL does not allow you to dump the mysql database.

- lock-all-tables: Common ApsaraDB RDS for MySQL users cannot use this option because they do not have the RELOAD privilege.

  Error message

  ```
  mysqldump: Couldn't execute 'FLUSH TABLES': Access denied; you need (at least one of) the RELOAD privilege(s) for this operation (1227)
  ```

- tab=dir_name: This option can be used only when mysqldump is run on the same physical machine as ApsaraDB RDS for MySQL instances. However, this option can be used together with the no-data option to obtain CREATE TABLE statements that create tables.

  ```
  # This option can be used with the no-data option to obtain the tab_name.sql file that contains the CREATE TABLE statement for each table in the test database.
  mysqldump  -no-defaults -uuser_name -ppass_word -hxxx.mysql.rds.aliyuncs.com -P3306 -set-gtid-purged=off -single-transaction -tab=/tmp -no-data test
  # The following error message appears if you use the no-data option to dump data:
  mysqldump  -no-defaults -uuser_name -ppass_word -hxxx.mysql.rds.aliyuncs.com -P3306 -set-gtid-purged=off -single-transaction -tab=/tmp test
  mysqldump: Got error: 1045: Access denied for user 'xxx'@'%' (using password: YES) when executing 'SELECT INTO OUTFILE'
  ```

4.4 How do I restore a downloaded RDS backup to a new RDS instance?

This topic describes how to restore data from a MySQL or SQL Server backup file to a new RDS instance.

Scenarios

This topic describes how to restore data to an apsaradb for RDS instance by using a backup client. You can restore data to a new RDS instance by using the backup client.

Note:

- Apsaradb RDS for PostgreSQL or PPAS does not support restoring data by using physical backup files to an apsaradb RDS for PPAS instance. You need to use a client to perform logical backup before releasing the instance. You need to use the pg_dump function to restore data to a new RDS instance when restoring data. For more information, see #unique_36.
- Currently, RDS for MariaDB does not allow you to download backup files.

Restore data in an apsaradb RDS for MySQL instance

You can recover data to the new RDS instance by following these steps:

1. Restore the backup data of an apsaradb for RDS instance to an on-premises database. For more information, see #unique_37 or #unique_38.
2. Use data transmission service (DTS) to migrate the local database to the new RDS instance. For more information, see #unique_39.

SQL Server Recovery method

You can restore data to a new RDS instance by migrating data from the OSS instance to the cloud. For more information, see #unique_40or #unique_41.

4.5 What do I do if the free quota of backup space is exceeded?

When the size of your backup files exceeds the free quota of backup space, you can increase the storage space or reduce backup space usage.

The backup files of an instance occupy backup space. Each ApsaraDB for RDS instance provides a quota of free backup space. When the space consumed by backup files exceeds this free quota, charges will be incurred.

Solutions

Use the following solutions to resolve the problem:

- Increase storage space (recommended): Because the free storage quota is 50% of the total storage space, increasing the storage space can increase the free quota. For more information about how to increase storage space, see the following topics:
  - Change the configuration of an ApsaraDB RDS for MySQL instance
  - Change the configuration of an ApsaraDB RDS for SQL Server instance
  - Change the configuration of an ApsaraDB RDS for PostgreSQL instance
  - Change the configuration of an ApsaraDB RDS for PPAS instance
  - Change the configuration of an ApsaraDB RDS for MariaDB TX instance
- Reduce the number of days for retaining backup data: You can set the data and log backup retention period to seven days. Backup files stored for more than seven days

will be automatically deleted to reduce storage space. For more information, see the following topics:

- #unique_48
- #unique_49
- #unique_50
- #unique_51
- #unique_52

Note:
The log backup retention period must be shorter than or equal to the data backup retention period.

- Reduce the backup cycle: You can specify to back up data two to three times a week. Perform backup at least twice in a week to guarantee data security. For more information, see the following topics:
  - #unique_48
  - #unique_49
  - #unique_50
  - #unique_51

  Note:
  By default, ApsaraDB RDS for MariaDB TX data is backed up on a daily basis and cannot be modified.
- Disable log backup: If you are sure that you do not need log backup, you can disable log backup. For more information, see the following topics:

  Warning:
  If you disable log backup, all log backup files will be deleted and data cannot be restored to previous points in time. Proceed with caution.

  - #unique_48
  - #unique_49
  - #unique_50
  - #unique_51
  - #unique_52

FAQ

Can uploading binlogs reduce backup space usage?

Uploaded binlogs are log backup files and thus still occupy backup space. In this case, backup space usage is not reduced.

5 Read-only Instances and Read/Write Splitting

5.1 Does ApsaraDB for RDS support read-only instances and read/write splitting?

Read-only instance

You can add read-only instances to your primary instance to process more read requests. For more information, see the following topics:

- #unique_55
- #unique_56
- #unique_57
- #unique_58

Read/write splitting

Read/write splitting is available to ApsaraDB for RDS instances that run MySQL and SQL Server. After you add read-only instances to your primary instance, you only need to enable read/write splitting. A read/write splitting endpoint is automatically generated to centrally receive and distribute query requests. For more information, see the following topics:

- Introduction to read/write splitting with ApsaraDB RDS for MySQL
- Introduction to read/write splitting with ApsaraDB RDS for SQL Server

Note:

ApsaraDB for RDS uses a primary/secondary architecture. In this architecture, the secondary instance does not process read or write requests but only serves as a backup to ensure high availability.

6 Performance

6.1 What occupies the capacity of new RDS for MySQL instances?

In RDS for MySQL instances, system files ib_logfile0 and ib_logfile1 occupy certain storage capacity.

After creating an RDS for MySQL instance, you can see that a few GB of storage space has been used. This is because of the system files ib_logfile0 and ib_logfile1.

The two log files are used to store the transaction log of the InnoDB engine table. Their size is always approximately 2 GB and cannot be changed. Due to the large size of the two files, the transaction log files do not need to be switched frequently when there are highly concurrent transactions. Therefore, the instance performance is improved.

7 Accounts

8 Databases

8.1 How do I perform DDL operations online on ApsaraDB RDS for MySQL?

Online DDL is a new feature available to ApsaraDB for RDS instances that run MySQL 5.6.

This feature allows you to perform data definition language (DDL) operations such as index creation on tables without blocking data manipulation language (DML) operations and SELECT queries that run concurrently with the DDL operations.

Note:

After you update your database engine version from an earlier version such as MySQL 5.5 to MySQL 5.6, you may not be able to perform DDL operations on a table for the first time because the table format is an earlier version. In this case, run the following command to convert the table format:

```
alter table <The name of the table whose format you want to convert> engine=innodb;
```

For information about more features supported by ApsaraDB RDS for MySQL, see #unique_66.

Limits

DDL operation	In-place supported	Table-copy required	DML concurrency permitted	Query concurrency permitted	Remarks
Create a common index on a table	Yes	No	Yes	Yes	None
Create a full-text index on a table	Yes	No	No	Yes	You must create the first full-text index by using the table-copy method. Then, you can create other full-text indexes on the table by using the in-place method.

DDL operation	In-place supported	Table-copy required	DML concurrency permitted	Query concurrency permitted	Remarks
Delete an index from a table	Yes	No	Yes	Yes	This DDL operation only modifies the metadata of the table.
Optimize a table	Yes	Yes	Yes	Yes	If a full-text index is created on the table, you cannot set the ALGORITHM option to INPLACE.
Set the default value in a column of a table	Yes	No	Yes	Yes	This DDL operation only modifies the metadata of the table.
Set the default value in an auto-increment column of a table.	Yes	No	Yes	Yes	This DDL operation only modifies the metadata of the table.
Add a foreign key constraint to a table	Yes	No	Yes	Yes	You can run the `set foreign_key_checks=0;` command to disable the foreign_key_checks option, so you do not need to copy the table.
Delete a foreign key constraint from a table	Yes	No	Yes	Yes	You can enable or disable the foreign_key_checks option.

DDL operation	In-place supported	Table-copy required	DML concurrency permitted	Query concurrency permitted	Remarks
Rename a column of a table	Yes	No	Yes	Yes	If this DDL operation only changes the column name without changing the data type, concurrent DML operations are permitted.
Add a column to a table	Yes	Yes	Yes	Yes	If the column you want to add is an auto_increment column, you cannot perform concurrent DML operations. You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.
Delete a column from a table	Yes	Yes	Yes	Yes	You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.
Change the sequence of columns in a table	Yes	Yes	Yes	Yes	You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.
Modify the Row_Format attribute of a table	Yes	Yes	Yes	Yes	You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.

DDL operation	In-place supported	Table-copy required	DML concurrency permitted	Query concurrency permitted	Remarks
Modify the Key_Block_Size attribute of a table	Yes	Yes	Yes	Yes	You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.
Set the values in a column of a table to NULL	Yes	Yes	Yes	Yes	You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.
Set the values in a column of a table to NOT NULL	Yes	Yes	Yes	Yes	This DDL operation only succeeds when the SQL_MODE option is set to STRICT_ALL_TABLES or STRICT_TRANS_TABLES and the column does not contain NULL values. You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.
Modify the data type of a column in a table	No	Yes	No	Yes	None

DDL operation	In-place supported	Table-copy required	DML concurrency permitted	Query concurrency permitted	Remarks
Add a primary key to a table	Yes	Yes	Yes	Yes	You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead. In addition, if you set the values in the column to NOT NULL, you cannot set the ALGORITHM option to INPLACE.
Delete the primary key of a table and add a new one	Yes	Yes	Yes	Yes	You can only set the ALGORITHM option to INPLACE when you execute one ALTER TABLE statement to delete the existing primary key and add a new one. You can set the ALGORITHM option to INPLACE, but this setting reorganizes the table data and increases the overhead.
Delete the primary key of a table	No	Yes	No	Yes	None
Convert the character set of a table	No	Yes	No	Yes	If the new character set uses a different encoding format, you must rebuild the table.

DDL operation	In-place supported	Table-copy required	DML concurrency permitted	Query concurrency permitted	Remarks
Specify a new character set for a table	No	Yes	No	Yes	If the new character set uses a different encoding format, you must rebuild the table.
Rebuild a table with the force option	Yes	Yes	Yes	Yes	If a full-text index is created on the table, you cannot set the ALGORITHM option to INPLACE.
Rebuild a table alter table ... engine= innodb	Yes	Yes	Yes	Yes	If a full-text index is created on the table, you cannot set the ALGORITHM option to INPLACE.
Set the persistent statistics attribute of a table	Yes	No	Yes	Yes	This DDL operation only modifies the metadata of the table.

- In-place supported: controlled by the ALGORITHM option in the DDL operation. This method consumes less disk space and I/O resources than the copy-table method.
- Copy-table required: controlled by the ALGORITHM option in the DDL operation. This method consumes more disk space and I/O resources than the in-place method.
- DML concurrency permitted: controlled by the LOCK option in the DDL operation.
- Query concurrency permitted: In most cases, queries concurrent with the DDL operation are permitted.
- For more information, see Online DDL Operations.

- The DDL operation on a table modifies the table metadata. Therefore, you may wait for the metadata lock on that table. For information about how to handle metadata locks, see #unique_67.
- The in-place method works in opposite to the table-copy method. However, even if the in-place method is specified, a DDL operation may still involve table copying, for example, the DDL operation used to add a column to a table.

Configuration suggestions

- ALGORITHM: When you perform a DDL operation, we recommend that you set this option to INPLACE to avoid performance deterioration caused by disk space usage or I/O issues. If the DDL operation does not support this setting, it returns an error.

  ```
  alter table area_bak algorithm=inplace, modify father text;

  ERROR 1846 (0A000): ALGORITHM=INPLACE is not supported. Reason: Cannot change column type INPLACE. Try ALGORITHM=COPY.
  ```

- LOCK: When you perform a DDL operation, we recommend that you set this option to NONE. This ensures the smooth execution of DML operations concurrent with the DDL operation. If the DDL operation does not support this setting, it returns an error.

  ```
  alter table area ALGORITHM=copy, lock=none,CONVERT TO CHARACTER SET utf8mb4;

  ERROR 1846 (0A000): LOCK=NONE is not supported. Reason: COPY algorithm requires a lock. Try LOCK=SHARED.
  ```

By default, ApsaraDB RDS for MySQL sets the ALGORITHM option to INPLACE and the LOCK option to NONE, so you do not need to configure these two options. However, if you are worried that a DDL operation may increase system loads or block DML operations on the target table, we recommend that you set the ALGORITHM option to INPLACE and the LOCK option to NONE separately for testing. Therefore, an error is returned when one of the two settings is not supported.

Example:

```
alter table area algorithm=inplace, lock=none, add index idx_fa (father);
```

Note:

We recommend that you perform all DDL operations during off-peak hours to avoid interruptions to your business.

If your database engine version (for example, MySQL 5.5) does not support online DDL operations, you can use the pt-online-schema-change tool of Percona.

For information about the syntax for ALTER TABLE, see ALTER TABLE Syntax.

Troubleshooting

When you perform an online DDL operation on a large table with concurrent DML operations, the following error may be returned:

```
alter table rd_order_rec add index idx_cr_time_detail (cr_time,detail);

ERROR 1799(HY000): Creating index 'idx_cr_time_detail' required more than 'innodb_onl
ine_alter_log_max_size' bytes of modification log. Please try again.
```

Cause

The system records any concurrent DML operations you perform during a DDL operation by creating a temporary log file when a table is modified or when an index is created. You can enlarge the temporary log file from the size specified by the innodb_sort_buffer_size parameter to the size specified by the nnodb_online_alter_log_max_size parameter.

If the temporary log file exceeds the maximum size, the DDL operation returns a failure message and rolls back all not committed concurrent DML operations. Therefore, we recommend that you use the innodb_online_alter_log_max_size parameter to specify a proper file size, so more concurrent DML operations are permitted. However, a larger file size increases the time taken by the DDL operation to lock the table and apply logged data to the table.

Solution

If your ApsaraDB for RDS instance runs MySQL 5.6 or 5.7, you can reconfigure the innodb_online_alter_log_max_size parameter in the ApsaraDB for RDS console. For more information, see #unique_68.

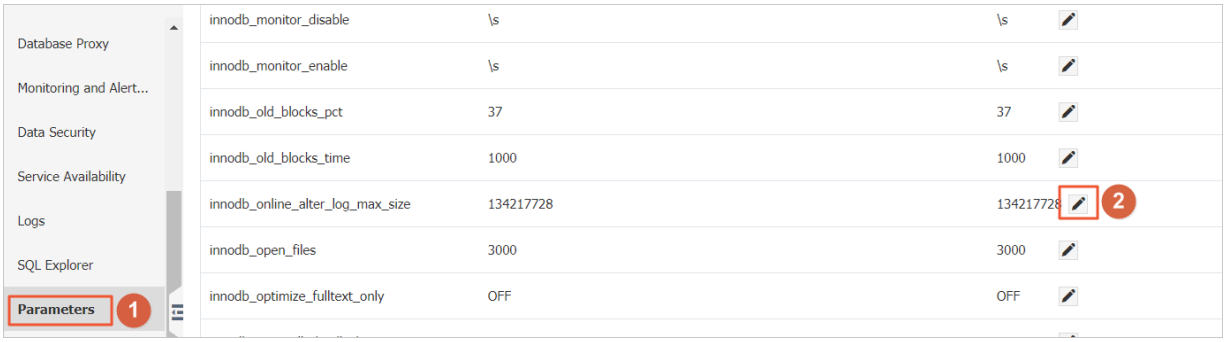

8.2 Limits on the size of a single table in RDS MySQL

The largest valid size of a single table is only constrained by the largest file size allowed by the operating system.

The storage capacity of an RDS MySQL instance can reach up to 2 TB. Therefore, the largest size of a single table is slightly less than 2 TB due to overheads of metadata. If an RDS MySQL instance has more than one table, the total size of these tables cannot exceed 2 TB.

If you have more questions, contact Alibaba Cloud technical support.

9 Logs

9.1 FAQ on binlogs

1. Q: Two binlog files have a similar start time and end time. Why has this occurred? Is the start time of one file consecutive with that of another?

   A: The two files contain binlogs that were generated during the backup of the primary and secondary nodes. As a result, the two files have a similar start time and end time. The start time of each file is not the same.

2. Q: Are binlog files compressed?

   A: Binlog files are not compressed.

3. Q: How are the generation, upload, and clearing of binlogs triggered?

   A: When more than 500 MB of log data is written to a binlog file, a new binlog file is generated. Backup logs are uploaded to OSS based on local backup settings, and then local logs are cleared.

10 SQL

10.1 What do I do if my ApsaraDB RDS for MySQL instance reports Error 1709?

Symptoms

When you create an index, the system may report the following error:

```
ERROR 1709 (HY000): Index column size too large. The maximum column size is 767 bytes.
```

Causes

The maximum indexed field length allowed by the InnoDB storage engine is 767 bytes. If you create an index on a multi-byte field or multiple fields, this error may be reported.

Note:

For example, utf8mb4 is a four-byte character set. The default indexed field length allowed by the utf8mb4 character set is 191 characters based on the following formula: 767 bytes/4 bytes per character ≈ 191 characters. If you create an index on a field of the varchar (255) or char (255) type, the index creation fails. For more information, see the MySQL documentation at the official MySQL website.

Solutions

1. Log on to the ApsaraDB for RDS console and navigate to the **Parameters** page. On the Editable Parameters tab, find the **innodb_large_prefix** parameter, change its value to ON or 1, and click **Apply Changes**.

Note:

After you set the **innodb_large_prefix** parameter to ON or 1, the maximum indexed field length allowed by InnoDB in the DYNAMIC or COMPRESSED row format increases to 3,072 bytes.

2. Specify the DYNAMIC or COMPRESSED row format during table creation. Example:

```
create table idx_length_test_02
(
  id int auto_increment primary key,
  name varchar(255)
)
ROW_FORMAT=DYNAMIC default charset utf8mb4;
```

Note:

For a created table, you can run the following commands to change its row format to DYNAMIC or COMPRESSED:

```
alter table <The name of the table> row_format=dynamic;
alter table <The name of the table> row_format=compressed;
```

11 Parameter